

UNITED STATES DISTRICT COURT
WESTERN DISTRICT OF TEXAS
WACO DIVISION

)
Intellectual Ventures I LLC and)
Intellectual Ventures II LLC,) Civil Action No. 6:19-cv-449
)
Plaintiffs,)
)
v.)
)
VMware, Inc.)
)
Defendant.) **JURY TRIAL DEMANDED**
)

COMPLAINT FOR PATENT INFRINGEMENT

Plaintiffs, Intellectual Ventures I LLC and Intellectual Ventures II LLC (together “IV”), for their complaint against defendant, VMware, Inc. (“VMware”), hereby alleges as follows:

THE PARTIES

1. Intellectual Ventures I LLC (“Intellectual Ventures I”) is a Delaware limited liability company having its principal place of business located at 3150 139th Avenue SE, Bellevue, Washington 98005.

2. Intellectual Ventures II LLC (“Intellectual Ventures II”) is a Delaware limited liability company having its principal place of business located at 3150 139th Avenue SE, Bellevue, Washington 98005.

3. Upon information and belief, VMware is a Delaware corporation with its headquarters located at 3401 Hillview Avenue, Palo Alto, California. VMware has regular and established places of business in this District, including two in Austin, Texas with over 700

employees. VMware also has at least two other offices in Texas—one in Coppell, Texas, and the other in Farmers Branch, Texas. Defendant may be served with process through its registered agent, Corporation Service Company, d/b/a CSC-Lawyers Incorporating Service Company, 211 E. 7th Street, Ste. 620, Austin, Texas 78701-3136

JURISDICTION

4. This is a civil action for patent infringement under the patent laws of the United States, 35 U.S.C. § 271, *et seq.* This Court has subject matter jurisdiction under 28 U.S.C. §§ 1331 and 1338(a).

5. This Court has general personal jurisdiction over VMware because VMware is engaged in substantial and not isolated activity at its regular and established places of business within this judicial district. This Court has specific jurisdiction over VMware because VMware has committed acts of infringement giving rise to this action and has established more than minimum contacts within this judicial district, such that the exercise of jurisdiction over VMware in this Court would not offend traditional notions of fair play and substantial justice.

6. Venue is proper in this judicial district pursuant to 28 U.S.C. §§ 1391(b)-(c) and 1400(b) because VMware maintains regular and established places of business and has committed acts of patent infringement within this judicial district.

FACTUAL BACKGROUND

7. Intellectual Ventures Management, LLC (“Intellectual Ventures”) was founded in 2000. Since then, Intellectual Ventures has been involved in the business of inventing. Intellectual Ventures creates inventions and files patent applications for those inventions; collaborates with others to develop and patent inventions; and acquires and licenses patents from individual inventors, universities, corporations, and other institutions. A significant aspect of

Intellectual Ventures' business is managing the plaintiffs in this case, Intellectual Ventures I and Intellectual Ventures II.

8. To create its own inventions, Intellectual Ventures has a staff of scientists and engineers who develop ideas in a broad range of fields, including agriculture, computer hardware, life sciences, medical devices, semiconductors, and software. Intellectual Ventures has invested millions of dollars developing such ideas and has filed hundreds of patent applications on its inventions every year, making it one of the top patent filers in the world. Intellectual Ventures has also invested in laboratory facilities to assist with the development and testing of new ideas.

9. Furthermore, Intellectual Ventures develops inventions by collaborating with inventors and research institutions around the world. For example, Intellectual Ventures has collaborated on inventions by selecting a technical challenge, requesting proposals for inventions to solve the challenge from inventors and institutions, selecting the most promising ideas, rewarding the inventors and institutions for their contributions, and filing patent applications on the ideas. Intellectual Ventures has invested millions of dollars in this way and has created a network of more than 4,000 inventors worldwide.

10. The founder of Intellectual Ventures is Nathan Myrhold, who worked at Microsoft from 1986 until 2000 in a variety of executive positions, culminating in his appointment as the company's first Chief Technology Officer ("CTO") in 1996. While at Microsoft, Mr. Myrhold founded Microsoft Research in 1991, and was one of the world's foremost software experts. Between 1986 and 2000, Microsoft became the world's largest technology company.

11. Under Mr. Myrhold's leadership, IV has acquired more than 70,000 patents covering some of the most important inventions of the Internet era. Many of these inventions coincided with Mr. Myrhold's successful tenure at Microsoft.

12. Two significant accomplishments of the Internet era are the related technologies of cloud computing and virtualization. Cloud computing enables ubiquitous access to shared pools of configurable computing system resources, such as memory, and higher-level services, such as virtual desktops or applications that run on these system resources. These resources and services, often collectively referred to as “the cloud,” can be rapidly provisioned with minimal management effort, often over the Internet. Virtualization allows for the creation of multiple software environments of dedicated resources that each simulate for its users a computing environment that was traditionally physically distinct such as a server or an operating system.

13. Three of the many beneficial consequences of cloud computing and virtualization are the consolidation of the server resources that businesses require to run a varied suite of software applications, the increased continuity of those server resources, and the outsourcing of those server resources to companies that operate large server clouds. Server resources are consolidated by replacing large numbers of physically distinct computer servers with a smaller number of computer servers, that each can host many virtual servers (also known as virtual machines). Each virtual machine (or “VM”) can execute the functions previously delivered by an entire computer server. Many VMs can run on a single physical server, which results in server consolidation and promotes scalability. Increased continuity is achieved in large part by enabling VMs to seamlessly move from one computer server to another computer server, for example, in the event that the first server becomes overloaded. Outsourcing of server resources requires technology that provides several different customers (i.e., users) access to a pool of those resources across a network, while providing each of those customers fault tolerance, performance and security isolation from the other customers accessing that same pool of resources.

14. The consolidation of servers by replacing a physical server with a virtual server has enabled services to migrate into giant centralized “clouds,” such as the Google or Amazon cloud, and more generally allows access to shared pools of configurable computing resources and services over the Internet. This enables services to be delivered to a given customer using centralized servers that seem as though they were located just a few feet away from the user, even if they are actually located hundreds or even thousands of miles away (hence the terms “virtual service” and “virtual server”).

15. The concept of transferring VMs, in the event of a computer server overload, to a different computer server has made VMs far less prone to outages, which in turn has changed how companies that provide virtualization technologies such as VMware are perceived. Such virtualization technology companies were previously associated with the limited function of server consolidation and thought of as relatively unreliable. With the more resilient architecture afforded by VMs that can move from physical server to physical server, however, virtualization companies are now also recognized as critical enablers of business continuity and availability.

16. Outsourcing of server resources to pools of physical servers, such as those found in data centers, enables large-scale and inexpensive delivery of virtual server services to a wide range of customers scattered across the Internet. The availability of large-scale and inexpensive virtual server services to the public at large has further enabled customers to focus less on building and maintaining their own network of physical computer servers and more on running their core businesses. It has also allowed small and mid-size businesses to enjoy the sorts of cutting edge computing technologies that were previously reserved for larger companies that could afford to build and maintain their own sophisticated networks of physical servers.

17. VMware provides virtualization solutions and services to its customers. VMware's product offerings include: vSphere, VMware's foundational virtualization platform; NSX, VMware's network virtualization solution; and vSAN, VMware's storage virtualization product. With 117 worldwide offices, VMware markets and sells these solutions and services throughout the globe, including in the United States and Texas.

THE PATENTS-IN-SUIT

The General Magic Patent

18. On May 24, 2011, the Patent and Trademark Office (PTO) issued United States Patent No. 7,949,752 ("the '752 patent"), titled NETWORK SYSTEM EXTENSIBLE BY USERS. The '752 patent is valid and enforceable. A copy of the '752 patent is attached as Exhibit A.

19. Intellectual Ventures I is the owner and assignee of all rights, title and interest in and to the '752 patent and holds all substantial rights therein, including the right to grant licenses, to exclude others, and to enforce and recover past damages for infringement of the '752 patent.

20. The inventions of the '752 patent were conceived at a famous 1990s technology company, General Magic, Inc. ("General Magic"). Founded in 1990 by members of the original Apple Macintosh development team, General Magic was a forerunner in the field of cloud computing and filed the original application for the '752 patent. General Magic's engineering team is viewed as one of the most talented in Silicon Valley history, and the fountain from which much of today's smartphone and online communication and commerce technology sprang.¹ The General

¹ See, eg.,: <https://www.cnet.com/news/at-premiere-of-general-magic-doc-tech-icons-consider-the-future/>; <https://medium.com/@writerjudy7/general-magic-award-winning-documentary-of-one-the-most-influential-tech-companies-youve-never-77d1c5b881aa> etc.

Magic team was even the subject of a recently released award-winning documentary. The inventors of the ‘752 patent were integral members of that team.

21. The term “cloud computing” was popularized to a large extent by General Magic. Computer industry trade publications cite to General Magic’s work as the earliest mention of cloud computing: “The term *cloud* was used to refer to platforms for distributed computing as early as 1993, when Apple spin-off General Magic and AT&T used it in describing their (paired) Telescript and PersonaLink technologies” (*see, e.g.*, <https://www.wired.com/2014/05/tech-time-warp-cloud-is-born/>).

22. The General Magic patent relates to a system, method and/or apparatus for running parts of computer programs, specifically data and instructions called “objects,” that operate in what is today called the cloud. The computer programs covered by the patent operate in a manner that enables the delivery from the cloud of highly reliable cloud and virtualization services to customers whose local computing platforms can be simple enough to support only a web browser.

The Ensim Patents

23. On December 31, 2013, the PTO issued United States Patent No. RE 44,686 (“the RE ‘686 patent”), titled DYNAMICALLY MODIFYING THE RESOURCES OF A VIRTUAL SERVER. The RE ‘686 patent is valid and enforceable. A copy of the RE ‘686 patent is attached as Exhibit B.

24. Intellectual Ventures I is the owner and assignee of all rights, title and interest in and to the RE ‘686 patent and holds all substantial rights therein, including the right to grant licenses, to exclude others, and to enforce and recover past damages for infringement of the RE ‘686 patent.

25. On September 20, 2011, the PTO issued United States Patent No. RE 42,726 (“the RE ‘726 patent”), titled DYNAMICALLY MODIFYING THE RESOURCES OF A VIRTUAL SERVER. The RE ‘726 patent is valid and enforceable. A copy of the RE ‘726 patent is attached as Exhibit C.

26. Intellectual Ventures I is the owner and assignee of all right, title and interest in and to the RE ‘726 patent and holds all substantial rights therein, including the right to grant licenses, to exclude others, and to enforce and recover past damages for infringement of the RE ‘726 patent.

27. The inventions of the RE ‘686 and RE ‘726 patents were conceived while the inventors worked at Ensim Corporation (“Ensim”). One of the inventors, Srinivasan Keshav, held the Cisco Systems and Canada Research chairs at the prestigious University of Waterloo in Canada.

28. The Ensim patents identified above relate to a system, method and/or apparatus for transferring VMs from one physical server to another physical server to help avoid outages, and thereby enable increased service continuity for business and consumer customers.

29. On December 27, 2011, the PTO issued United States Patent No. RE 43,051 (“the RE ‘051 patent”), titled ENABLING A SERVICE PROVIDER TO PROVIDE INTRANET SERVICES. The RE ‘051 patent is valid and enforceable. A copy of the RE ‘051 patent is attached as Exhibit D.

30. Intellectual Ventures I is the owner and assignee of all rights, title and interest in and to the RE ‘051 patent and holds all substantial rights therein, including the right to grant licenses, to exclude others, and to enforce and recover past damages for infringement of the RE ‘051 patent.

31. The inventions claimed in the RE ‘051 patent were conceived while the inventors worked at Ensim Corporation. The inventors, Pawan Goyal and Peter Newman, are both highly respected in their field with decades of experience at companies such as Ensim, Microsoft, IBM, Symantec and Nokia, to name a few. Dr. Goyal holds a Ph.D. in computer science from the University of Texas at Austin, and Dr. Newman received his Ph.D. in packet switching for integrated services from the University of Cambridge. Both men have published extensively in industry journals such as IEEE Communications Magazine and are named inventors on numerous foreign and U.S. patents. Dr. Goyal is a co-inventor, for example, of key software container technologies, which are one of the most common ways of packaging and delivering software in today’s cloud-driven landscape.

32. The RE ‘051 patent relates to a system, method and/or apparatus for networked communications with a group of virtual servers that each belong to different customers, but that all run on a common pool of physical servers. The communications covered by the RE ‘051 patent use “network tunnels” to better isolate a customer’s virtual server from other customers, faults and security vulnerabilities that may appear elsewhere in the common pool of physical servers.

The 3Leaf Systems Patent

33. On March 25, 2014, the PTO issued United States Patent No. RE 44,818 (“the RE ‘818 patent”), titled QUALITY OF SERVICE IN VIRTUAL COMPUTING ENVIRONMENTS. The RE ‘818 patent is valid and enforceable. A copy of the RE ‘818 patent is attached as Exhibit E.

34. Intellectual Ventures II is the owner and assignee of all rights, title and interest in and to the RE ‘818 patent and holds all substantial rights therein, including the right to grant

licenses, to exclude others, and to enforce and recover past damages for infringement of the RE ‘818 patent.

35. The inventions claimed in the RE ‘818 patent were conceived while the inventors worked at 3Leaf Systems, Inc., a pioneer in the network virtualization field. The company received over \$67M in funding, partnering with industry giants like IBM and Intel to develop network virtualization solutions that by one account are “changing the way storage and server virtualization are done.” 3Leaf was granted numerous patents, including the RE ‘818 patent, whose inventors went on to work in network virtualization at companies like Google.

36. The RE ‘818 patent relates to a system, method and/or apparatus for the operation of virtualized input/output (I/O) systems in virtual server environments. The recent adoption of virtualized I/O subsystems has increased the flexibility with which I/O resources can be assigned to physical and virtual servers, and has reduced the cost of operating and maintaining these same I/O resources. Virtualized I/O subsystems have also, however, increased the demands placed on the network that underpins each virtual server environment. That network must now handle a greater number of communications between the underlying I/O subsystems and physical servers (“I/O communications”), which must be moved across the network with increasingly varied and exacting networking requirements. The RE ‘818 helps address such increased requirements by allocating network access to a given I/O communication in an increasingly sophisticated manner, based in part on the physical/virtual network interfaces involved with the I/O communication, and the type of operation being supported by the I/O communication.

COUNT I
(VMware’s Infringement of U.S. Patent No. 7,949,752)

37. Paragraphs 1-36 are reincorporated by reference as if fully set forth herein.

38. The elements claimed by the ‘752 patent, taken alone or in combination, were not well-understood, routine or conventional to one of ordinary skill in the art at the time of the invention. Rather, the ‘752 patent claims and teaches, *inter alia*, an improved way to deliver a network-based user-customized service to end-users accessing the service using standard web browsers or other client applications capable of interpreting markup language code, which were not present in the state of the art at the time of the invention. The invention improved upon existing services delivered over a network, which were difficult to customize or extend, by augmenting the programs at the servers that deliver those services, with a server-based agent system, and monitoring resource usage of that agent system to protect other users and service providers from being deprived of their resources.

39. Instead of the services being directly provided to a remote client by a conventional server-side service program such as voicemail services that ran on telephone networks in the 1990s, or mainframe service provided to IBM display-terminal devices as was practiced in the prior art, the services in the present invention are provided to the user via an agent system residing on the server that is invocable using a client application such as a browser that can, for example, interpret markup language such as HTML. The agent system then invokes and uses the server programs and resources on behalf of its user to deliver the services to the user’s client. With this capability, an agent’s use of the server programs and resources can be customized or extended using specialized application program interfaces that exist between the server programs and resources on one hand, and the agents on the other hand. Thus, the resulting web-based services provided to the user are more extensible and customizable relative to prior art systems that required specialized modifications to server-based programs to satisfy the custom requirements of particular subscribers.

40. Compared to the prior art, the claimed system for delivering web-based services is also more resilient against excessive resource usage by users, since the resource allocations being requested by users through each agent can be more easily and continuously monitored/enforced due in large part to the presence of the network-based agent.

41. The invention represented a technical solution to an unsolved technological problem. The written description of the '752 patent describes, in technical detail, each of the limitations in the claims, allowing a person of skill in the art to understand what those limitations cover, and therefore what was claimed, and also understand how the non-conventional and non-generic ordered combination of the elements of the claims differ markedly from what had been conventional or generic in the industry at the time of the inventions of the '752 patent. More specifically, the claims of the '752 patent each recite a network-based (e.g., server side) agent. The agent (VM) can be invoked upon the system receiving a URL defining a type of event (e.g., instantiating a VM) and identifying the agent (VM). In so doing, a service (e.g., a Windows operating system or application) can be used by the agent (VM) on behalf of a user for performing an operation, and server resources are consumed by the agent as it performs said operation.

42. The system covered by the asserted claims, therefore, differs markedly from the conventional and generic systems in use at the time of this invention, which *inter alia* lacked the claimed combination of network-based agents, agent servers, and programmability. Embodiments of the present invention further include a service wrapper associated with the service, the service wrapper for cooperating with the agent server and mediating interactions between the service and the agent.

43. As described above, the '752 patent is drawn to solving a specific, technical problem arising in the context of network-based services. Consistent with the problem addressed

being rooted in such network-based computing environments, the ‘752 patent’s solutions naturally are also rooted in that same technology and cannot be performed with pen and paper or in the human mind.

44. VMware has directly infringed, and continues to directly infringe, literally and/or by the doctrine of equivalents, individually and/or jointly, at least claims 1-4, 6, 9-11, 13-14, and 22-26 of the ‘752 patent by making, using, testing, selling, offering for sale and/or importing into the United States products and/or services covered by one or more claims of the ‘752 patent. VMware’s products and/or services that infringe the ‘752 patent include, but are not limited to, Horizon View, vSphere, ESX/ESXi, vRealize Operations Manager, vCenter Server, Connection Server, and any other VMware products and/or services, either alone or in combination, that operate in substantially the same manner.

45. Claim 9 of the ‘752 patent is reproduced below:

9. *A system for performing user customized network-based operations, comprising:*

- a processor; and*
- a memory storing instructions, execution of which by the processor causes the system to perform operations comprising:*

 - receiving data for creating a network-based agent;*
 - invoking, in response to receiving a URL defining a type of event and identifying the network-based agent, an execution of the network-based agent wherein the execution of the network-based agent comprises using a service and a service resource configured to be consumed by the network-based agent when the network-based agent performs the operation, and wherein an amount of the service resource is exhausted upon being consumed by the network-based agent, and*
 - communicating a result of the operation over a network communications link.*

46. As one non-limiting example, VMware Horizon in combination with vSphere includes a system for performing user customized network-based operations as claimed in the ‘752 patent. For example, the Horizon products and/or services can provide users with remote access to virtual operating systems, desktops and applications (service(s)) hosted on a vSphere backend.

Aspects of the system such as the mode of delivery and the configuration settings for virtual applications and desktops can be made customizable as is explained in further detail below:

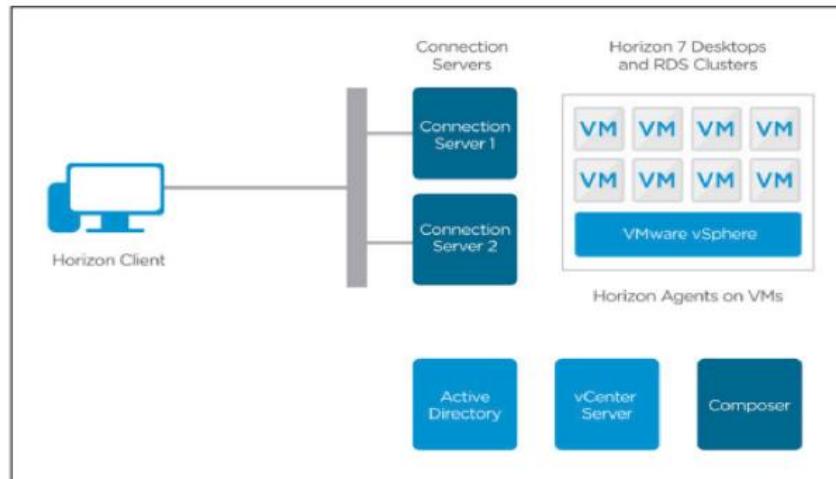
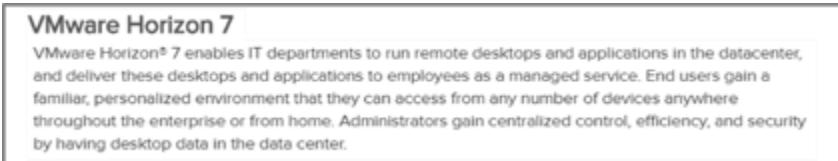
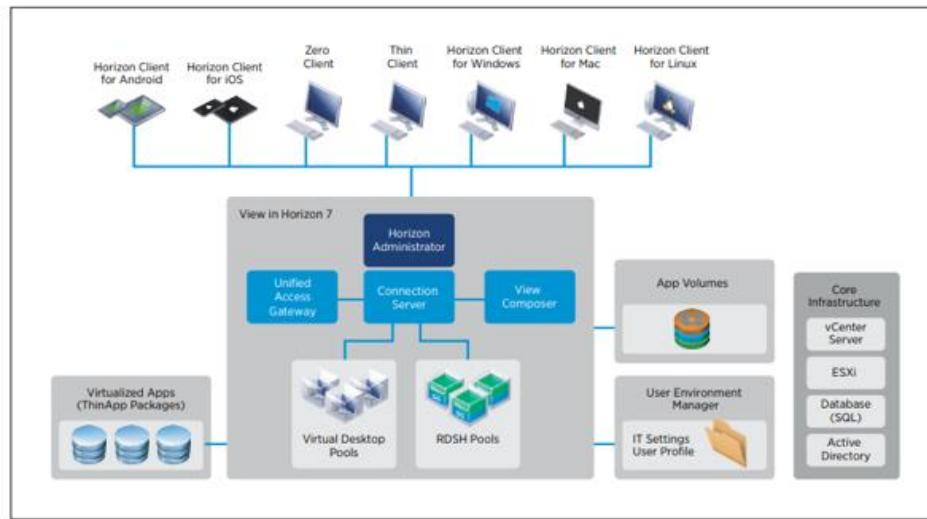


Figure 16: View Logical Architecture

Architecture and Components

View contains key components and integrated products that work together.



47. An example of a typical deployment of Horizon includes at least one network-based agent in the form of an ESX/ESXi-based virtual machine (e.g. a virtual desktop or virtual

application) managed by the vSphere backend. Each virtual machine (VM) is instantiated on an underlying physical host machine that is managed by a combination of vSphere and ESX/ESXi. The VM is allocated at least one virtual processor (vCPU) and at least some virtual memory (vRAM, vMemory), which are each derived from the host machine's underlying physical resource (physical CPU and physical RAM/memory).

48. The vCPU and vRAM/vMemory on the VM are used to execute and store program instructions as if they were physical resources running on a physical machine. Thus, Horizon in conjunction with vSphere and ESX/ESXi, includes a processor and a memory for storing executable program instructions that allow the system to perform useful operations:

ESXi Hardware Requirements

To install or upgrade ESXi, your hardware and system resources must meet the following requirements:

- ESXi 6.7 requires a host machine with at least two CPU cores.
- ESXi 6.7 supports 64-bit x86 processors released after September 2006. This includes a broad range of multi-core processors. For a complete list of supported processors, see the VMware compatibility guide at <http://www.vmware.com/resources/compatibility>.
- ESXi 6.7 requires a minimum of 4 GB of physical RAM. It is recommended to provide at least 8 GB of RAM to run virtual machines in typical production environments.

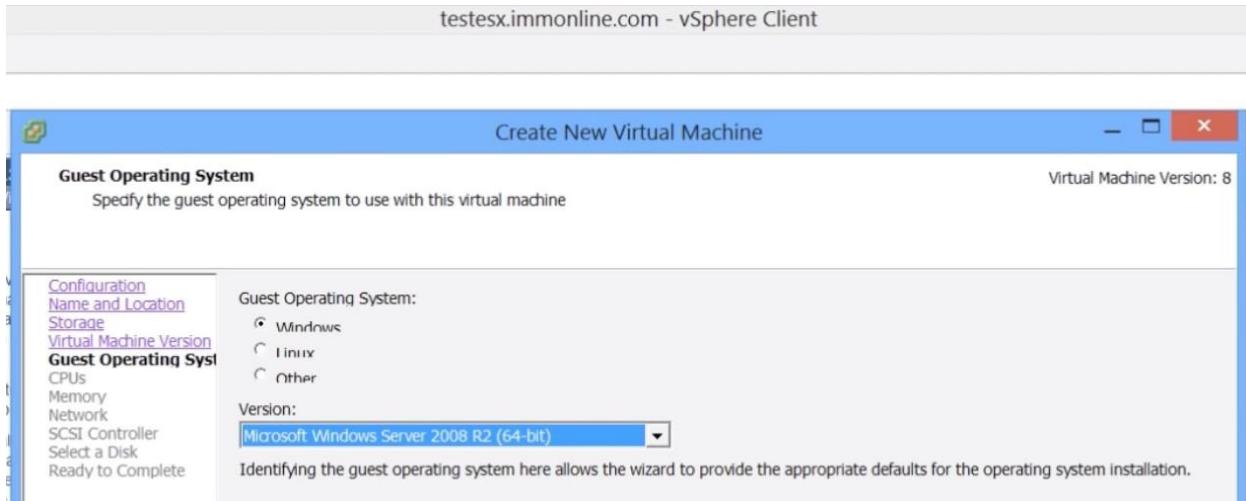
Resource Consumers

Virtual machines are resource consumers.

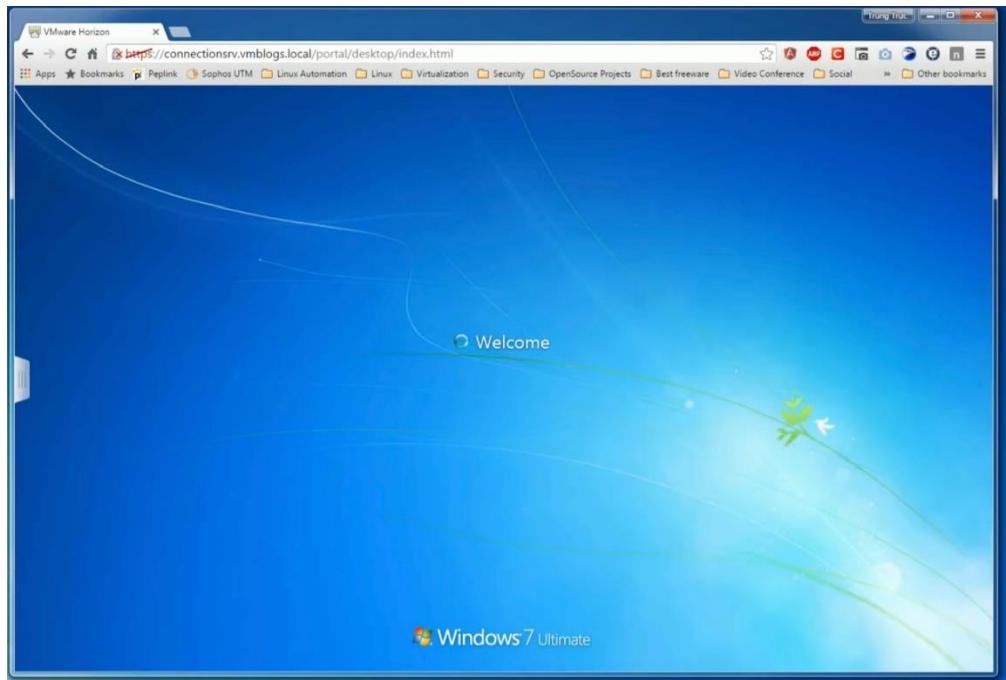
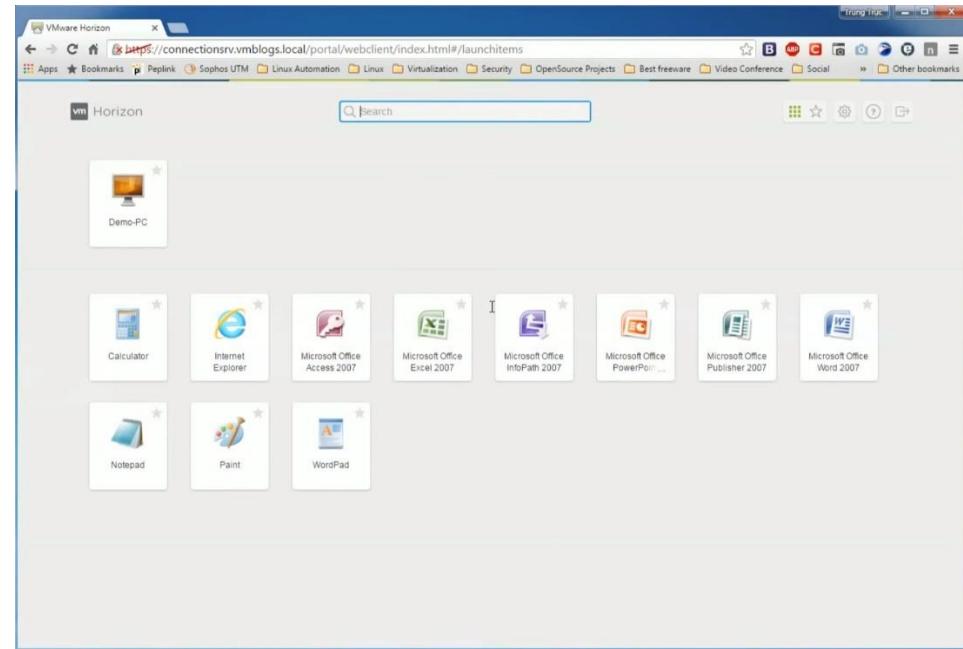
The default resource settings assigned during creation work well for most machines. You can later edit the virtual machine settings to allocate a share-based percentage of the total CPU, memory, and storage I/O of the resource provider or a guaranteed reservation of CPU and memory. When you power on that virtual machine, the server checks whether enough unreserved resources are available and allows power on only if there are enough resources. This process is called admission control.

49. Furthermore, continuing with the above example, Horizon, in conjunction with at least vSphere and ESX/ESXi, enables a deployment for receiving data for creating a network-

based agent. For example, a Horizon deployment provides a platform on which VMs (agents) can be created and hosted on a vSphere backend (e.g., agent server). A vSphere client (e.g. vSphere Web Client) can be used to create/modify/delete (receiving data) the VM in a Horizon deployment:

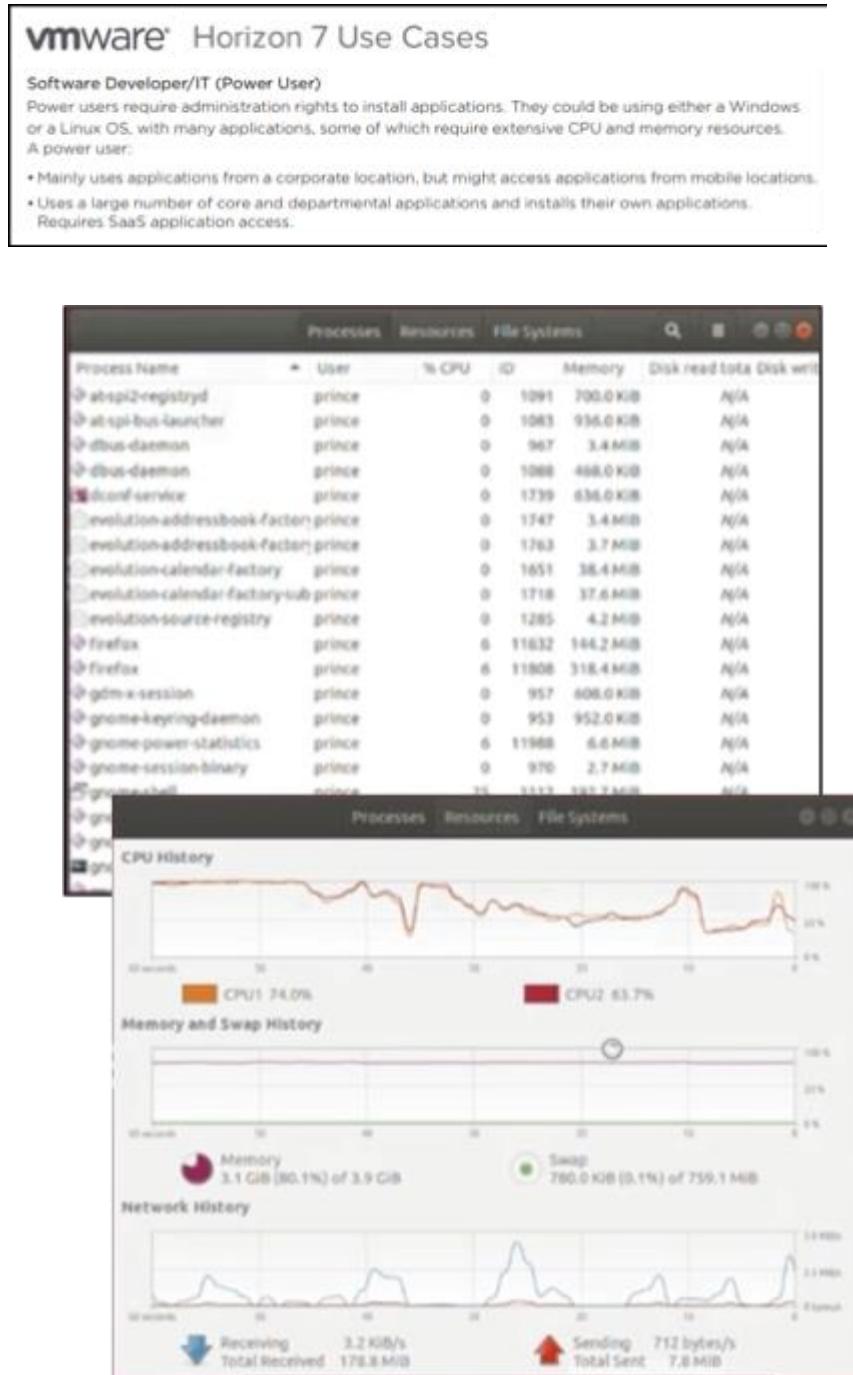


50. Horizon further provides for invoking, in response to receiving a URL defining a type of event and identifying the network-based agent, execution of the network-based agent. For example, a launch URL is received at one or more servers, such as a View Connection Server or Direct Connect Plugin. The URL identifies the network-based agent (VM) and defines the type of event requested (e.g., the instantiation of a VM). In response, the VM is instantiated and made available for remote access, thus an execution of the network-based agent is invoked:



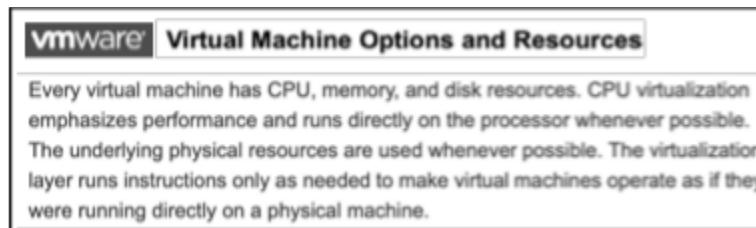
51. Horizon also includes the limitation wherein the execution of the network-based agent comprises using a service and a service resource configured to be consumed by the network-based agent when the network-based agent performs the operation. For example, a virtual desktop or application that has been invoked within a vSphere backend includes a plurality of processes

running thereon (services used). While in use, each service occupies some amount of vCPU and vRAM/vMemory allocated to the VM (service resources consumed):



52. Horizon also includes the limitation wherein an amount of the service resource is exhausted upon being consumed by the network-based agent. As discussed above, each VM is

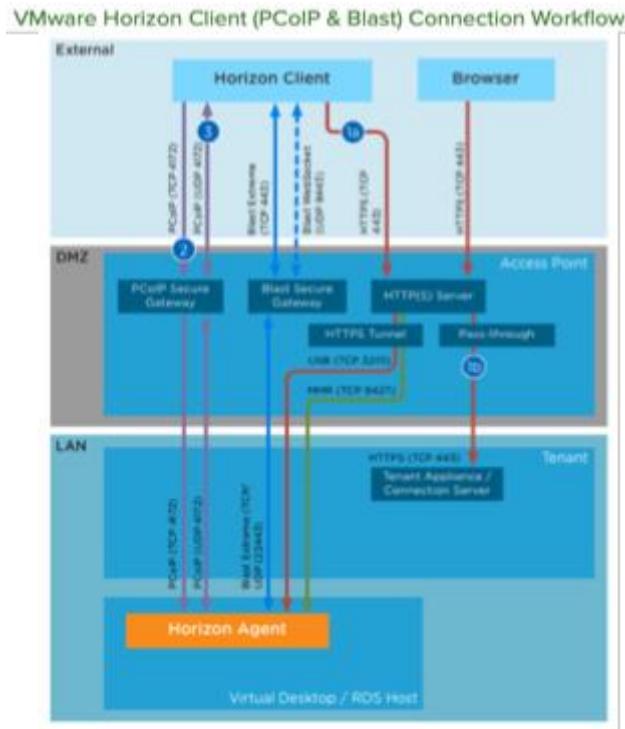
allocated by vSphere/ESXi a certain amount of vCPU and vRAM/vMemory (service resources), which correspond to a portion of the physical CPU and RAM/memory resources available on a host machine. These physical resources are finite and, therefore, when service resources are in use, the underlying physical resources are occupied and cannot be used by other processes until they are released by the VM (service resources are exhausted):



<u>Linux VM (Running on ESXi Host)</u>	<u>ESXi Host</u>
<p>Summary Monitor Configure Permissions Data</p> <p>Guest OS: Ubuntu Linux (64-bit) Compatibility: ESXi 6.5 and later (VM) VMware Tools: Running, version:10304 DNS Name: vcenter-linux IP Addresses: 192.168.171.73 Host: 192.168.171.25</p> <p>Launch Web Console Launch Remote Console</p>	<p>192.168.171.25 ACTIONS</p> <p>Summary Monitor Configure Permissions VMS Datastore</p> <p>Hypervisor: VMware ESXi 6.5.0_8294253 Model: HP ProDesk 600 G3 SFF Processor Type: Intel(R) Core(TM) i5-7500 CPU @ 3.40GHz Logical Processors: 4 NICs: 1 Virtual Machines: 1 State: Connected Uptime: 1 days</p>
<p>CPU History Processes Resources File systems</p> <p>CPU History: CPU1 31.0% (orange), CPU2 31.4% (red)</p> <p>Memory and Swap History: Memory 3.1 GB (79.2%) of 3.9 GB, Swap 780.0 kB (0.1%) of 759.1 MB</p>	<p>CPU: Free: 6.43 GHz Used: 7.2 GHz Capacity: 13.83 GHz</p> <p>Memory: Free: 257.14 MB Used: 3.63 GB Capacity: 3.88 GB</p>

53. Horizon also provides for communicating a result of the operation over a network communications link. For instance, the VM computes the results of virtual desktop or application launch activity (invoking an execution of the network-based agent) and the results of subsequent VM process activity (using a service) and communicates (in conjunction with vSphere) these

results to a remote client-facing user interface (e.g. Horizon desktop client or Horizon HTML Access web client):



54. Additionally, VMware has been, and currently is, an active inducer of infringement of the ‘752 patent under 35 U.S.C. § 271(b) and contributory infringement of the ‘752 patent under 35 U.S.C. § 271(c) either literally and/or by the doctrine of equivalents.

55. VMware has actively induced, and continues to actively induce, infringement of the ‘752 patent by intending that others use, offer for sale, or sell in the United States, products and/or services covered by one or more claims of the ‘752 patent, including but not limited to VMware’s Horizon View, vSphere, ESX/ESXi, vRealize Operations Manager, vCenter Server, Connection Server, and any VMware product and/or service, alone or in combination, that operates in materially the same manner. VMware provides these products and/or services to others, such as customers, resellers and end-user customers, who, in turn, use, provision for use, offer for sale, or

sell in the United States products and/or services that directly infringe one or more claims of the '752 patent.

56. VMware has contributed to, and continues to contribute to, the infringement of the '752 patent by others by knowingly providing products and/or services that, when installed and configured result in a system as intended by VMware, that directly infringes one or more claims of the '752 patent.

57. VMware knew of the '752 patent, or should have known of the '752 patent, but was willfully blind to its existence. Upon information and belief, VMware has had actual knowledge of the '752 patent since at least as early as the service upon VMware of this Complaint. By the time of trial, VMware will have known and intended (since receiving such notice) that its continued actions would infringe and actively induce and contribute to the infringement of one or more claims of the '752 patent.

58. VMware has committed, and continues to commit, affirmative acts that cause infringement of one or more claims of the '752 patent with knowledge of the '752 patent and knowledge or willful blindness that the induced acts constitute infringement of one or more claims of the '752 patent. As an illustrative example only, VMware induces such acts of infringement by its affirmative actions of intentionally providing hardware and or software components that when used in their normal and customary way as desired and intended by VMware, infringe one or more claims of the '752 patent and/or by directly or indirectly providing instructions on how to use its products and/or services in a manner or configuration that infringes one or more claims of the '752 patent, including those found at one or more of the following:

- www.vmware.com/products/horizon.html
- www.vmware.com/products/esxi-and-esx.html

- www.vmware.com/products/vsphere.html
- www.vmware.com/pdf/view401_admin_guide.pdf
- docs.vmware.com/en/VMware-Horizon-7/index.html
- www.youtube.com/watch?v=3OvrKZYnzjM
- www.youtube.com/watch?v=oJMCXhZeHIA
- www.youtube.com/watch?v=EvXn2QiL3gs
- www.youtube.com/watch?v=xpYDM6sZOWY

59. VMware has also committed, and continues to commit, contributory infringement by, *inter alia*, knowingly selling products and/or services that when used cause the direct infringement of one or more claims of the ‘752 patent by a third party, and which have no substantial non-infringing uses, or include a separate and distinct component that is especially made or especially adapted for use in infringement of the ‘752 patent and is not a staple article or commodity of commerce suitable for substantial non-infringing use.

60. As a result of VMware’s acts of infringement, Plaintiffs have suffered and will continue to suffer damages in an amount to be proved at trial.

COUNT II

(VMWare’s Infringement of U.S. Patent No. RE 44,686)

61. Paragraphs 1-60 are reincorporated by reference as if fully set forth herein.

62. The elements claimed by the RE ‘686 patent, taken alone or in combination, were not well-understood, routine or conventional to one of ordinary skill in the art at the time of the invention of the RE ‘686 patent, around 2000. Rather, the patent teaches and claims an improved way to monitor and dynamically adjust resource usage by a VM that represented a novel, and non-obvious approach that was not present in the state of the art at that time. The invention improved upon existing VM technology by automatically moving operating VMs from a first computer

server to a second computer server if the VM needed additional resources that were not presently available on the first computer server but were available on the second computer server.

63. The invention represented a technical solution to an unsolved technological problem. The written description of the RE ‘686 patent describes, in technical detail, each of the limitations in the claims, allowing a person of skill in the art to understand what those limitations cover, and therefore what was claimed, and also understand how the non-conventional and non-generic ordered combination of the elements of the claims differ markedly from what had been conventional or generic in the industry at the time of the invention of the RE ‘686 patent.

64. More specifically, the claims of the RE ‘686 patent each require a first computer server for hosting a VM, an indication that the first computer server does not have the resources to accommodate a requested increase in resource allocation by the VM, a host transfer signal that indicates a second server has the resources to accommodate the requested increase in resource allocation, and the transferring of the VM from the first computer server to the second computer server. The systems covered by the asserted claims therefore differ markedly from the conventional and generic systems in use at the time of this invention, which *inter alia* did not involve automatically moving VMs that were running on overloaded computer servers to another computer server.

65. As described above, the RE ‘686 patent is drawn to solving a specific, technical problem arising in the context of managing resources for VMs in a networked computing environment. Consistent with the problem addressed being rooted in such VMs running in a networked computing environment, the RE ‘686 patent’s solutions are also rooted in that same technology and cannot be performed with pen and paper or in the human mind.

66. VMware has directly infringed, and continues to directly infringe, literally and/or by the doctrine of equivalents, individually and/or jointly, the RE ‘686 patent by making, using, testing, selling, offering for sale and/or importing into the United States products and/or services covered by at least claims 1-3 and 5-7 of the RE ‘686 patent. VMware’s products and/or services that infringe the RE ‘686 patent include, but are not limited to, vSphere, vCenter Server, vMotion, Distributed Resource Scheduler (“DRS”), VMware’s ESX/ESXi hypervisor, and any other VMware products and/or services, either alone or in combination, that operate in materially the same manner:

67. Claim 5 of the RE ‘686 patent is reproduced below:

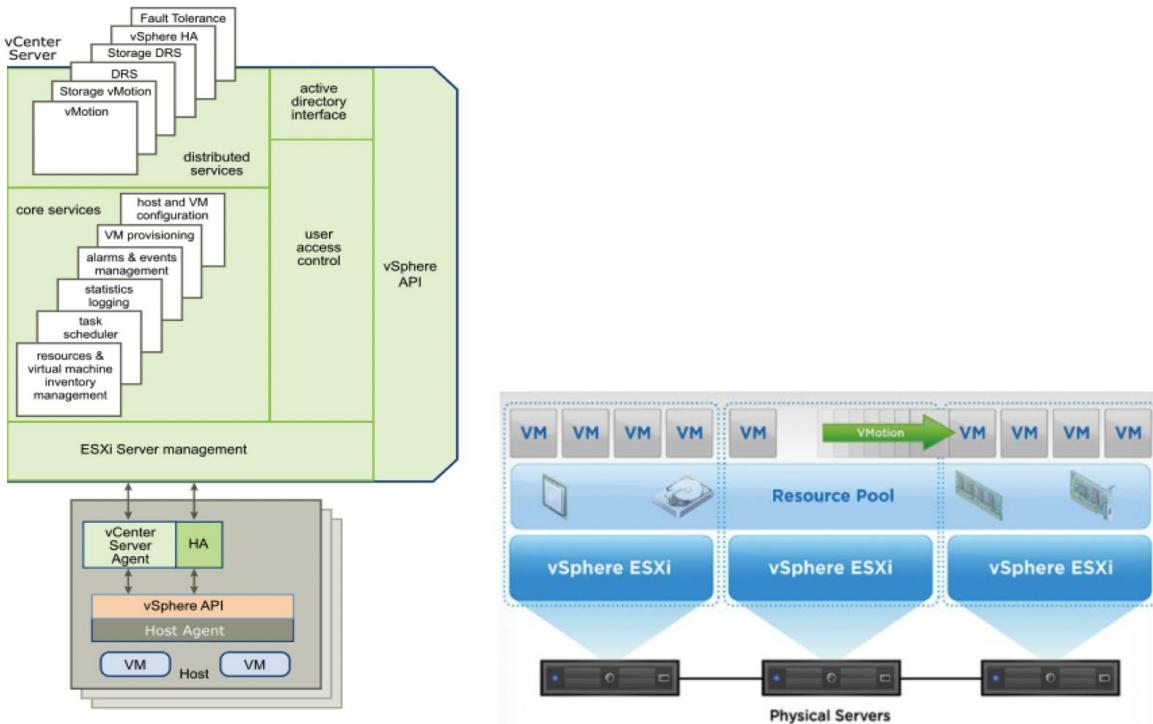
5. A method performed by a computing device, having a processor and memory, for modifying the computer resources allocated to a virtual server operating in a first physical host of multiple physical hosts, comprising:

receiving an indication that a first physical host is overloaded, wherein the indication is based on a determination that a virtual server is overloaded and wherein the determination that a virtual server is overloaded is based on one or more resource unavailable messages resulting from denied requested to modify a resource allocation;

determining that a second physical host can accommodate the requested modified resource allocation; and

generating a physical host transfer signal that indicates the second physical host and transferring the virtual server from the first physical host to the second physical host.

68. As one non-limiting example, a deployment of vSphere, vMotion, DRS, and ESX/ESXi provides for modifying the computer resources allocated to a VM operating in a first physical host of multiple physical hosts. For example, the vSphere platform can manage and monitor a deployment’s virtual infrastructure, which can include a plurality of physical host machines and the VMs hosted thereon, and in conjunction with vSphere’s DRS and vMotion functionality is capable of modifying allocated physical resources among the VMs.

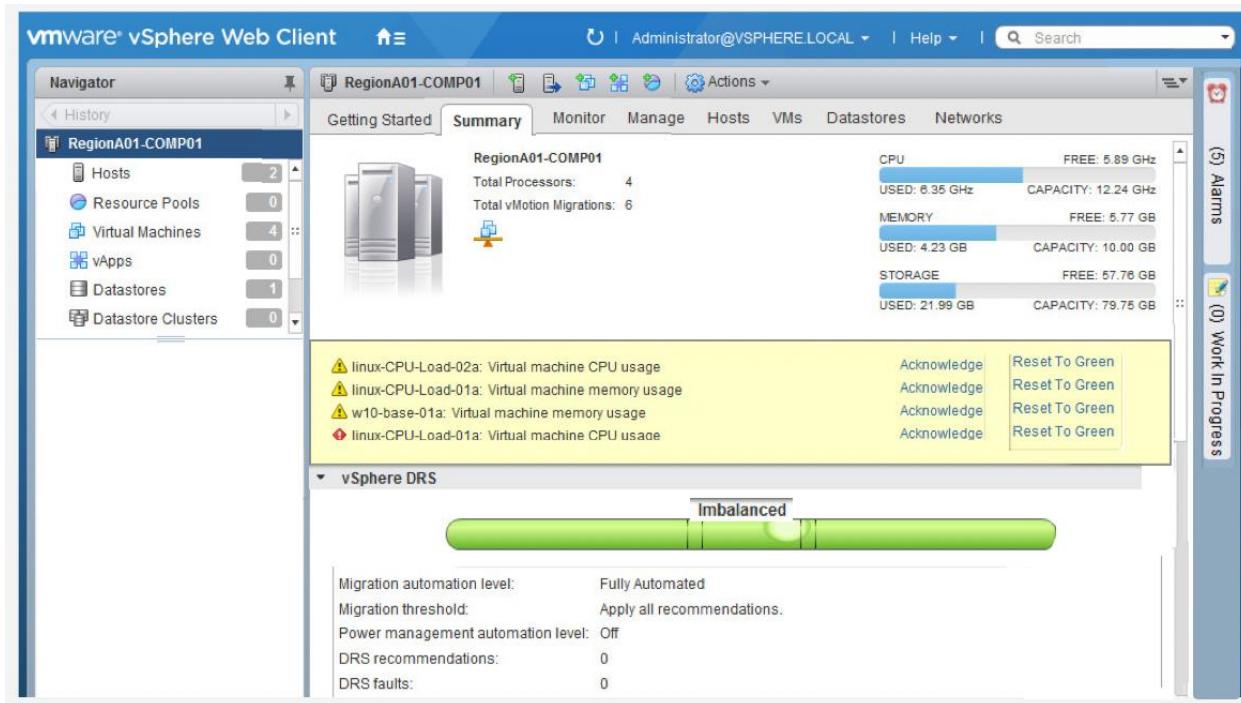


69. For example, the vSphere platform can receive an indication that a first physical host is overloaded, wherein the indication is based on a determination that a virtual server is overloaded and wherein the determination that a virtual server is overloaded is based on one or more resource unavailable messages resulting from denied requests to modify a resource allocation. The vSphere platform including vMotion and DRS can monitor the physical resources allocated to the deployment's VMs and their resource usage, including when the VMs are denied resources. The vSphere platform can further determine that a second physical host can accommodate the requested modified resource allocation, evaluate resource-usage of different hosts, and accordingly, can determine the destination host which can satisfy the resource needs of the VM:

How DRS Works

From time to time, VMs' workloads may change, and with many VMs with changing workloads, there can be imbalance in the cluster. Each of these can degrade application performance. DRS solves these problems by regularly monitoring the cluster balance state once every five minutes, by default, and then takes the necessary actions to fix any imbalance. DRS automatically determines which virtual machines would benefit from a move to another host and live migrates the VM onto the new host using vMotion. In this way, DRS ensures each virtual machine in the cluster gets the host resources—like memory and CPU—that it needs.

70. Continuing with the above example, the vSphere platform can generate a physical host transfer signal to transfer the VM from the first physical host to the second physical host if the first physical host is overloaded. For example, the vSphere platform and the vMotion/DRS functionality can generate a migration recommendation that indicates a second physical host and can result in the transfer of the VM from the first physical host to the second physical host if the first physical host is overloaded.



71. VMware and its customers operating, for example, vSphere environments with vMotion and DRS functionality directly infringe at least claims 5-7 of the '686 patent. In order to

obtain the benefit of the vSphere and ESX/ESXi functionality, including the vMotion and DRS features, VMware's customers must operate vSphere as required by VMware and therefore, if any action by VMware's customer is necessary for direct infringement of any claims (including at least claims 5-7), such actions are attributable to VMware.

72. Additionally, VMware has been, and currently is, an active inducer of infringement of the RE '686 patent under 35 U.S.C. § 271(b) and contributory infringer of the RE '686 patent under 35 U.S.C. § 271(c) either literally and/or by the doctrine of equivalents.

73. VMware has actively induced, and continues to actively induce, infringement of the RE '686 patent by intending that others use, offer for sale, or sell in the United States, products and/or services covered by one or more claims of the RE '686 patent, including but not limited to VMware's vSphere platform, ESX/ESXi hypervisor, vMotion and DRS functionality, and any VMware product and/or service that operates in materially the same manner. VMware provides these products and/or services to others, such as customers, resellers and end-user customers, who, in turn, use, provision for use, offer for sale, or sell in the United States products and/or services that directly infringe one or more asserted claims of the RE '686 patent.

74. VMware has contributed to, and continues to contribute to, the infringement of the RE '686 patent by others by knowingly providing products and/or services that, when installed and configured result in a system as intended by VMware, directly infringes one or more claims of the RE '686 patent.

75. VMware knew of the RE '686 patent, or should have known of the RE '686 patent, but was willfully blind to its existence. Upon information and belief, VMware has had actual knowledge of the RE '686 patent since at least as early as the service upon VMware of this Complaint. By the time of trial, VMware will have known and intended (since receiving such

notice) that its continued actions would infringe and actively induce and contribute to the infringement of one or more claims of the RE ‘686 patent.

76. VMware has committed, and continues to commit, affirmative acts that cause infringement of one or more claims of the RE ‘686 patent with knowledge of the RE ‘686 patent and knowledge or willful blindness that the induced acts constitute infringement of one or more claims of the RE ‘686 patent. As an illustrative example only, VMware induces such acts of infringement by its affirmative action by intentionally providing hardware and/or software components that when used in their normal and customary way as designed and intended by VMware, infringe one or more claims of the RE ‘686 patent and/or by directly or indirectly providing instructions on how to use its products and services in a manner or configuration that infringes one or more claims of the RE ‘686 patent, including those found at one or more of the following:

- <https://www.vmware.com/products/horizon.html>
- <https://www.vmware.com/products/horizon-cloud-virtual-desktops.html>
- <https://www.vmware.com/products/vcenter-server.html>
- <https://www.vmware.com/products/vsphere.html>
- <https://www.vmware.com/products/vsphere-hypervisor.html>
- <https://blogs.vmware.com/vsphere/2016/10/introducing-vsphere-6-5.html>
- https://www.vmware.com/pdf/vmware_drs_wp.pdf
- https://www.vmware.com/pdf/vmotion_datasheet.pdf
- <https://blogs.vmware.com/vsphere/2016/10/drs-migration-thresholds.html>
- <https://kb.vmware.com/s/article/1017926>
- <https://www.vmware.com/products/vsphere/vmotion.html>

77. VMware has also committed, and continues to commit, contributory infringement by, *inter alia*, knowingly selling products and/or services that when used cause the direct infringement of one or more claims of the RE ‘686 patent by a third party, and which have no substantial non-infringing uses, or include a separate and distinct component that is especially made or especially adapted for use in infringement of the RE ‘686 patent and is not a staple article or commodity of commerce suitable for substantial non-infringing use.

78. As a result of VMware’s acts of infringement, Plaintiffs have suffered and will continue to suffer damages in an amount to be proved at trial.

COUNT III
(VMware’s Infringement of U.S. Patent No. RE 42,726)

79. Paragraphs 1- 78 are reincorporated by reference as if fully set forth herein.

80. The elements claimed by the RE ‘726 patent, taken alone or in combination, were not well-understood, routine or conventional to one of ordinary skill in the art at the time of the invention of the RE ‘726 patent, around 2000. Rather, the patent teaches and claims an improved way to monitor and dynamically adjust resource usage by a VM that represented a novel, and not obvious approach that was not present in the state of the art at that time. The invention improved upon existing VM technology by incorporating VM monitoring, load balancing and automatically moving VMs from a first computer server to a second computer server if the VM needed additional resources that were not presently available on the first computer server but were available on the second computer server.

81. The invention represented a technical solution to an unsolved technological problem. The written description of the RE ‘726 patent describes, in technical detail, each of the limitations in the claims, allowing a person of skill in the art to understand what those limitations cover, and therefore what was claimed, and also understand how the non-conventional and non-

generic ordered combination of the elements of the claims differ markedly from what had been conventional or generic in the industry at the time of the invention of the RE ‘726 patent.

82. More specifically, the claims of the RE ‘726 patent each require one of a plurality of physical hosts for hosting a VM (virtual server), an VM resource monitor for monitoring resource denials to the VM and indicating that a VM is potentially overloaded, a load balancer for determining whether the VM is overloaded and indicating a host transfer is required, and functionality to move the VM from a first physical host to a second physical host if it is determined that the required resources are available on a second physical host. The systems covered by the asserted claims therefore differ markedly from the conventional generic systems in use at the time of this invention, which *inter alia* did not involve automatically moving operating VMs that were running on overloaded physical hosts to another physical host.

83. As described above, the RE ‘726 patent is drawn to solving a specific, technical problem arising in the context of managing resources for VMs in a networked computing environment. Consistent with the problem addressed being rooted in such VMs running in a networked computing environment, the RE ‘726 patent’s solutions are also rooted in that same technology and cannot be performed with pen and paper or in the human mind.

84. VMware has directly infringed, and continues to directly infringe, literally and/or by the doctrine of equivalents, individually and/or jointly, the RE ‘726 patent by making, using, testing, selling, offering for sale and/or importing into the United States products and/or services covered by at least claims 1-13 of the RE ‘726 patent. VMware’s products and/or services that infringe the RE ‘726 patent include, but are not limited to, vSphere, vMotion, Distributed Resource Scheduler (“DRS”), VMware’s ESX/ESXi hypervisor, and any VMware product and/or service that operates in materially the same manner.

85. Claim 1 of the RE ‘726 patent is reproduced below:

1. A network system for dynamically modifying the computer resources allocated to a virtual server, the network system comprising a plurality of physical hosts, the virtual server operating in a first physical host, the computer resources allocated to the virtual server being specified as a quality of service guarantee, the network system comprising:

a virtual server resource monitor communicatively coupled to the first physical host and configured to monitor resource denials and to send a virtual server overloaded signal in response to the resource denials;

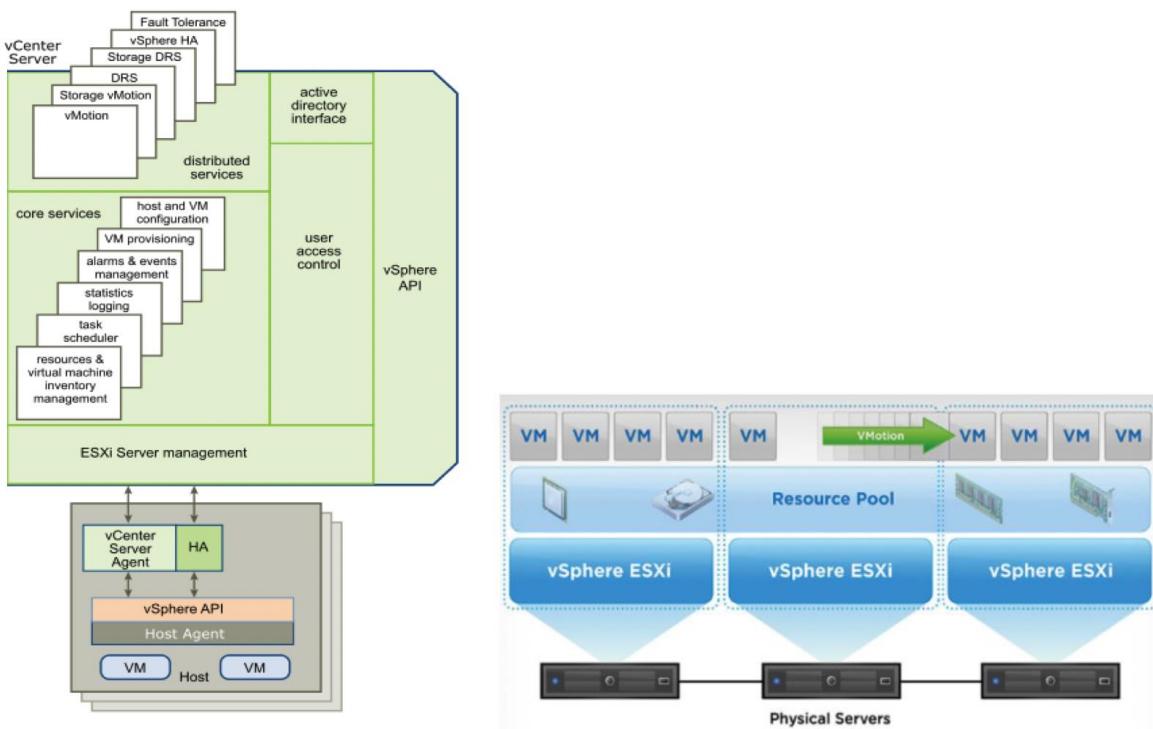
a virtual server resource modifier communicatively coupled to the first physical host and configured to receive the virtual server overloaded signal and, to modify a resource allocation for the virtual server and to send a virtual server resource modification signal;

a load balancing module communicatively coupled to the plurality of physical hosts and configured to receive the virtual server resource modification signal and to determine whether the first physical host is overloaded and, in response to a determination that the first physical host is overloaded, to send a physical host transfer signal that indicates a second physical host; and

a dynamic virtual server mover communicatively coupled to the plurality of physical hosts and configured to receive the physical host transfer signal, and in response to the physical host transfer signal, to transfer the virtual server from the first physical host to the second physical host.

86. As one non-limiting example, vSphere, vMotion, DRS, and ESX/ESXi infringe by providing a system for monitoring and modifying the computer resources allocated to a VM operating in a first physical host of multiple physical hosts as claimed in at least claims 1-13 of the RE ‘726 patent. For example, the vSphere platform can manage and monitor a site’s virtual infrastructure, which can include a plurality of physical host machines and the VMs hosted thereon, thereby performing each of the elements of at least claims 1-13. As a non-limiting example, VMware’s vSphere, and ESX/ESXi products include DRS and vMotion that perform the steps of at least claims 1-13. VMware’s customers operating, for example, vSphere environments with vMotion and DRS functionality directly infringe at least claims 1-13 of the ‘726 patent. In

order to obtain the benefit of the vSphere and ESX/ESXi functionality, including the vMotion and DRS features, VMware's customers must operate vSphere as required by VMware and therefore, if any action by VMware's customer is necessary for direct infringement of any claims (including at least claims 1-13), such actions are attributable to VMware.

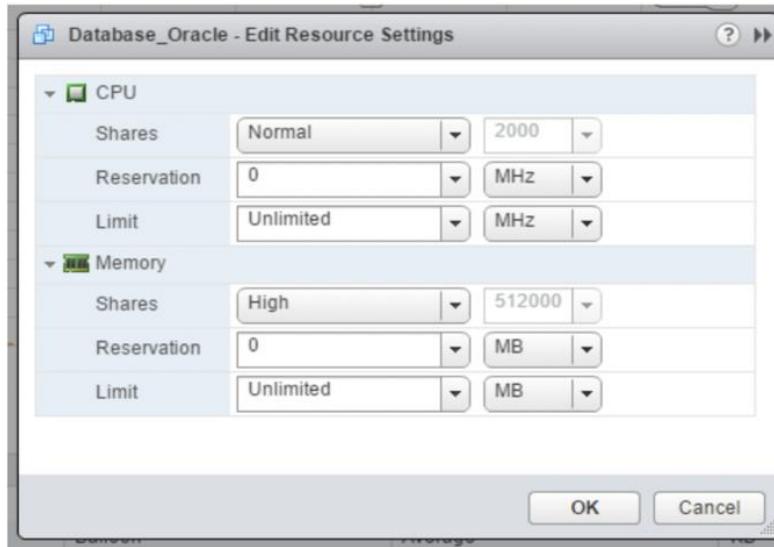


DRS computes a resource entitlement for each virtual machine, based on virtual machine and resource pool configured shares, reservations, and limits settings, as well as the current demands of the virtual machines and resource pools. What the virtual machine demands is not necessarily what it deserves or is entitled to because of the virtual machine and resource pool configurations.

Shares

Shares provide you a way to prioritize resources for VMs when there is competition in the cluster. They can be set at a VM or a resource pool level.

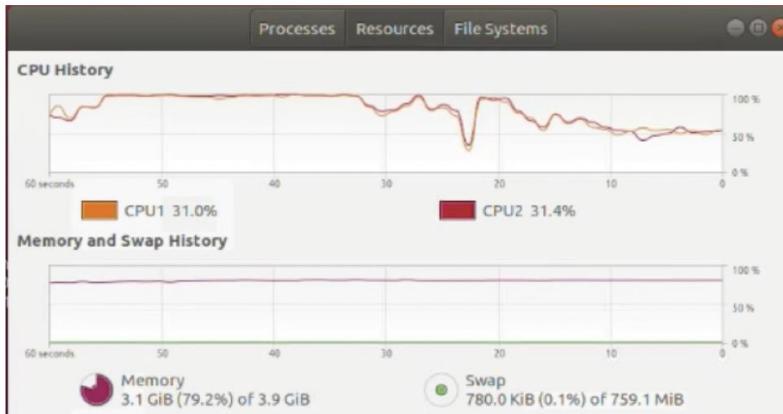
By default, a cluster has a resource pool hierarchy, with the root resource pool (the cluster itself) at the top, and all VMs as its children. Shares are defined as numbers for all the sibling VMs under this root resource pool. Shares are distributed equally, by default, on a per-resource basis (per-vCPU and per-unit of memory). This means that by default, a VM with more configured resources will get more shares than a VM with fewer resources. During resource contention, resources available at the root resource pool are shared among the children based on their shares' values.



87. For example, the vSphere platform (including vMotion and DRS) can monitor physical hosts on which VMs are instantiated and the resources used by such VMs, including any resource denials. Furthermore, the vSphere platform can send a VM overloaded signal in response to resource denials and modify a VM resource allocation for a VM in response to the overload signal. The vSphere platform also comprises a load balancer which can receive a resource modification signal in response to modification of a VM's resource allocations and determine whether a first physical host is overloaded. In response to a determination that a first physical host is overloaded the vSphere platform can send a physical host transfer signal to indicate a second physical host capable of providing the denied resources and transferring the VM from the first to the second physical host as a result:

The vCenter Server agent acts as a small vCenter Server to perform the following functions:

- Relays and enforces resource allocation decisions made in vCenter Server, including those that the DRS engine sends.
- Passes virtual machine provisioning and configuration change commands to the host agent.
- Passes host configuration change commands to the host agent.
- Collects performance statistics, alarms, and error conditions from the host agent and sends them to the vCenter Server.



88. VMware and its customers operating, for example, vSphere environments with vMotion and DRS functionality directly infringe at least claims 1-13 of the RE '726 patent. In order to obtain the benefit of the vSphere and ESX/ESXi functionality, including the vMotion and DRS features, VMware's customers must operate vSphere as required by VMware and therefore, if any action by VMware's customer is necessary for direct infringement of any claims (including at least claims 1-13), such actions are attributable to VMware.

89. Additionally, VMware has been, and currently is, an active inducer of infringement of the RE '726 patent under 35 U.S.C. § 271(b) and contributory infringer of the RE '726 patent under 35 U.S.C. § 271(c) either literally and/or by the doctrine of equivalents.

90. VMware has actively induced, and continues to actively induce, infringement of the RE '726 patent by intending that others use, offer for sale, or sell in the United States, products and/or services covered by one or more claims of the RE '726 patent, including but not limited to VMware's vSphere platform, ESX/ESXi hypervisor, vMotion and DRS functionality, and any VMware product and/or service that operates in materially the same manner. VMware provides these products and/or services to others, such as customers, resellers and end-user customers, who, in turn, use, provision for use, offer for sale, or sell in the United States products and/or services that directly infringe one or more asserted claims of the RE '726 patent.

91. VMware has contributed to, and continues to contribute to, the infringement of the RE ‘726 patent by others by knowingly providing products and/or services that when installed and configured result in a system as intended by VMware, that directly infringes one or more claims of the RE ‘726 patent.

92. VMware knew of the RE ‘726 patent, or should have known of the RE ‘726 patent, but was willfully blind to its existence. Upon information and belief, VMware has had actual knowledge of the RE ‘726 patent since at least as early as the service upon VMware of this Complaint. By the time of trial, VMware will have known and intended (since receiving such notice) that its continued actions would infringe and actively induce and contribute to the infringement of one or more claims of the RE ‘726 patent.

93. VMware has committed, and continues to commit, affirmative acts that cause infringement of one or more claims of the RE ‘726 patent with knowledge of the RE ‘726 patent and knowledge or willful blindness that the induced acts constitute infringement of one or more claims of the RE ‘726 patent. As an illustrative example only, VMware induces such acts of infringement by its affirmative action by intentionally providing hardware and/or software components that when used in their normal and customary way as designed and intended by VMware, infringe one or more claims of the RE ‘726 patent and/or by directly or indirectly providing instructions on how to use its products and services in a manner or configuration that infringes one or more claims of the RE ‘726 patent, including those found at one or more of the following:

- <https://www.vmware.com/products/horizon.html>
- <https://www.vmware.com/products/horizon-cloud-virtual-desktops.html>
- <https://www.vmware.com/products/vcenter-server.html>

- <https://www.vmware.com/products/vsphere.html>
- <https://www.vmware.com/products/vsphere-hypervisor.html>
- <https://blogs.vmware.com/vsphere/2016/10/introducing-vsphere-6-5.html>
- https://www.vmware.com/pdf/vmware_drs_wp.pdf
- https://www.vmware.com/pdf/vmotion_datasheet.pdf
- <https://blogs.vmware.com/vsphere/2016/10/drs-migration-thresholds.html>
- <https://kb.vmware.com/s/article/1017926>
- <https://www.vmware.com/products/vsphere/vmotion.html>

94. VMware has also committed, and continues to commit, contributory infringement by, *inter alia*, knowingly selling products and/or services that when used cause the direct infringement of one or more claims of the RE ‘726 patent by a third party, and which have no substantial non-infringing uses, or include a separate and distinct component that is especially made or especially adapted for use in infringement of the RE ‘726 patent and is not a staple article or commodity of commerce suitable for substantial non-infringing use.

95. As a result of VMware’s acts of infringement, Plaintiffs have suffered and will continue to suffer damages in an amount to be proved at trial.

COUNT IV
(VMware’s Infringement of U.S. Patent No. RE 43,051)

96. Paragraphs 1-95 are reincorporated by reference as if fully set forth herein.

97. The elements claimed by the RE ‘051 patent, taken alone or in combination, were not well-understood, routine or conventional to one of ordinary skill in the art at the time of the inventions of the RE ‘051 patent, around 2000. The RE ‘051 patent improved upon existing remote networking services (such as virtual hosting solutions that did not use tunnels) by enabling a service provider to:

- provide a customer with a virtual private server running on a physical server at a service provider site such as a service provider data center, with its own private address space;
- deploy many such virtual servers on one or more physical servers at the service provider data center in a multitenant configuration such that multiple customers (“tenants”) are each provided access to one or more of said virtual servers with a private address space;
- facilitate the exchange of privately addressed transmissions between each customer and her/his virtual server using external tunnels that traverse the local/regional network(s) separating the client site from the service provider data center; and
- differentiate, at a gateway near the boundary of the service provider data center and the local/regional network(s), transmissions moving to/from different virtual servers within the service provider data center using internal tunnels that extend within the service provider data center’s network from the gateway to the customer’s virtual servers.

98. The foregoing combination of features helps the service provider provide a number of isolation benefits to its customers including ensuring that the data of each customer is inaccessible to other customers, that problems in one customer’s service do not compromise another customer’s service, and that customers get their own private IP addresses (which provides more network security for the customers). Said features also enable the service provider to readily scale its business by adding more physical servers and more virtual private servers without compromising the security or quality of the private networking service it offers customers. The RE ‘051 patent, therefore, provides customers with more secure, reliable and scalable remote networking services without requiring a dedicated physical network topology for each customer.

99. By contrast, some prior private virtual server services used a dedicated physical server at the service provider data center, to provide each customer with a private addressing

scheme that was secure and unique to that customer. This prior art system, however, resulted in costly overhead for service providers and limited the scalability of the system. Other prior private networking services used one physical server for several customers in providing a virtual hosting service, but were unable to provide customers with the RE ‘051 patent’s isolation advantages as listed above.

100. The inventors of the RE ‘051 patent realized that by deploying an end-to-end tunneling scheme in which tunnel endpoints extend within a service provider data center to the edge of a virtual server hosted on one of its physical servers, each customer could get a dedicated virtual server having an address space that is capable of overlapping the address spaces of other virtual servers (meaning each customer’s traffic could be isolated and granting customers flexibility to use whatever internal addressing scheme they desire, which makes the system more secure and reliable). The inventors also realized that deploying such internal tunnels enabled several virtual servers to be deployed on a single physical server without compromising security or scalability.

101. The RE ‘051 patent represented a technical solution to an unsolved technological problem. The written description of the RE ‘051 patent describes in technical detail each of the limitations in the claims, allowing a person of skill in the art to understand what those limitations cover, and therefore what was claimed, and also understand how the non-conventional and non-generic ordered combination of the elements of the claims differ markedly from what had been conventional or generic in the industry at the time of the inventions of the RE ‘051 patent. More specifically, as an example, claim 3 of the RE ‘051 patent recites a host physical server containing multiple virtual servers that each support overlapping private address spaces in providing private

network services to remote users, and a tunneling mechanism for connecting each customer to its respective virtual server.

102. The claim goes on to specify how this is achieved through a lookup mechanism that includes storing customer lookup and forwarding information which associates physical interfaces with tunnel identifiers. The claim also specifies using the forwarding information to identify physical interfaces and tunnel identifiers associated with the network address of customer's virtual server. This mechanism allows a given customer's traffic to remain isolated from other customers' traffic as it moves between the customer's originating client device and that customer's virtual server (i.e., using an overlapping private address space). The system covered by the asserted claims therefore differs significantly from the conventional and generic systems in use at the time of this invention, which among other things, lacked the ability for a customer to communicate with a remotely located virtual server that was using its own private address space, and that was deployed on a physical server with lots of other virtual servers.

103. As described above, the RE '051 patent is drawn to solving a specific technological problem arising in the context of providing secure remote private network services with overlapping private addresses via virtual servers and tunneling. Consistent with the problem addressed being rooted in such computer-based virtual private networking environments, the RE '051 patent's solutions naturally are also rooted in the same technology and cannot be performed with a pen and paper or in the human mind.

104. VMware has directly infringed, and continues to directly infringe, literally and/or by the doctrine of equivalents, individually and/or jointly, at least claims 1 and 3-6 of the RE '051 patent by making, using, testing, selling, offering for sale and/or importing into the United States products and/or services covered by one or more claims of the RE '051 patent. VMware's products

and/or services that infringe the RE ‘051 patent include, but are not limited to, NSX, NSX-T, vCloud Director, vSphere, ESXi, and any other VMware products and/or services, either alone or in combination, that operate in materially the same manner.

105. Claim 3 of the RE ‘051 patent is represented below:

3. In a system comprising a host computer containing multiple virtual servers that each support a private network address space wherein the private network address spaces of two or more of the virtual servers overlap, a method for providing private network services using private addresses in a location remote from private network users, the method comprising:

storing customer lookup information and customer forwarding information, the customer lookup information specifying associations between physical interfaces and tunnel identifiers identifying tunnels for private networks and multiple customer forwarding tables, the customer forwarding information associating network addresses with physical interfaces and tunnel identifiers;

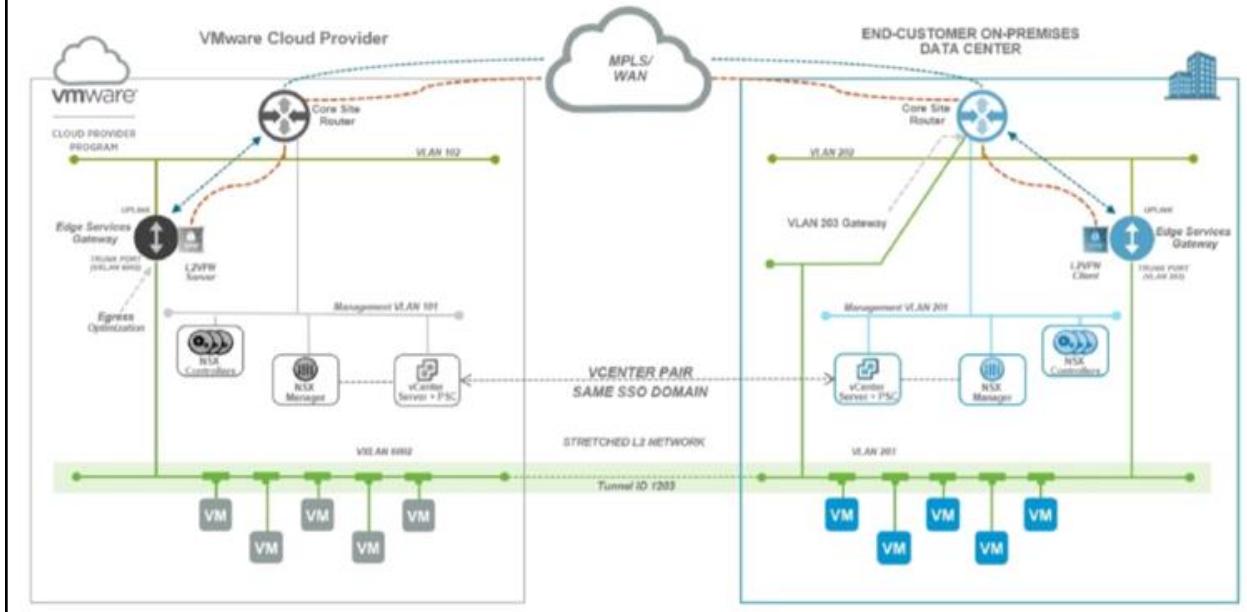
receiving, over a tunnel, a transmission on a physical interface having an interface identifier, the transmission identifying a tunnel identifier;

determining the correct customer forwarding information from the customer lookup information using the physical interface identifier and the tunnel identifier;

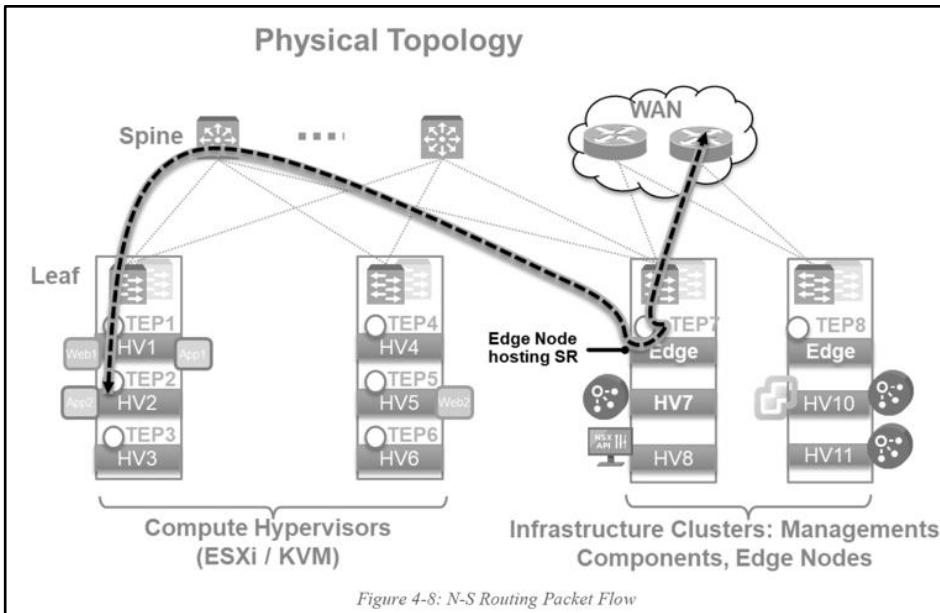
using the customer forwarding information to identify a physical interface and tunnel identifier associated with a network address of the transmission; and

sending the transmission to the network address on the identified physical interface using the identified tunnel identifier.

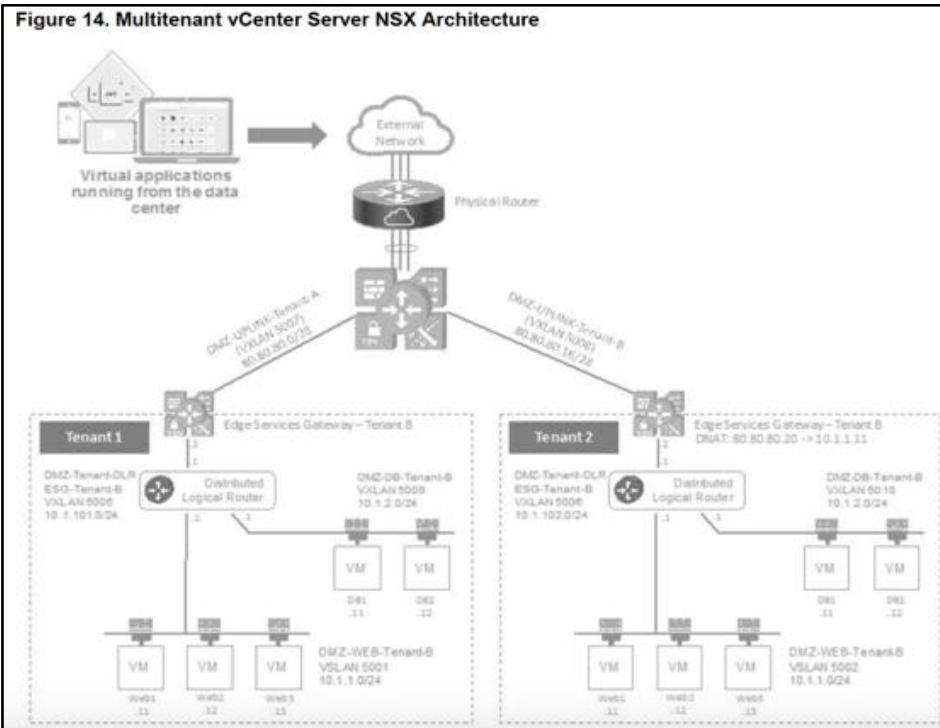
106. As one non-limiting example, VMware infringes by providing (and using) products and/or services that include a system comprising a host computer containing multiple virtual servers (e.g. compute VMs hosted on an ESXi hypervisor) that each support a private network address space wherein the private network addresses of virtual servers can overlap. For example, a VMware NSX deployment providing Layer 2 VPN (L2VPN) services to multiple tenants enables the provision of virtual servers with overlapping private network addresses such that each tenant’s network address space is isolated:

Figure 2. VMware NSX to VMware NSX Stretched L2VPN

Overlapping customer addresses are separated by the Org VDC networks whose VXLAN backing creates the same Layer 2 separation as traditional VLAN-backed networks. The only places at which the overlapping addresses could therefore clash is if the separated networks were then connected to shared routing devices. In the examples used throughout this document, each tenant's Org VDC has a dedicated Edge Services Gateway, and is connected to the respective customer WAN over a discrete vCloud Director external network. This VLAN-backed network typically terminates on a dedicated, per-tenant WAN CE router or a shared multi-tenant PE router in which the each VLAN is internally mapped to per-tenant VRF (as described in Section 4.2, vCloud Director Multitenant Data Center Networking in vSphere). The Edge Services Gateway, CE router, or PE VRF maintain independent routing tables, allowing each customer to use identical addresses within their tenant Organizations without affecting other tenants.



107. The exemplary NSX system described above can provide private network services using private addresses and virtual servers (as explained above) in a location remote from the user/customer (typically a service provider's data center):



108. The exemplary NSX system described above can store customer lookup information and customer forwarding information, the lookup information specifying associations between physical interfaces and tunnel identifiers, and the forwarding information associating network addresses with physical interfaces and tunnel identifiers. For example, the NSX control plane stores customer lookup information such as, client MAC or IP addresses, VLAN ID, VXLAN ID and port numbers which associate a customer with physical interfaces and tunnel identifiers. The customer forwarding information includes destination server MAC or IP address, VLAN ID, VXLAN ID, VTEP, and port number, which are associated with a destination addresses, physical interfaces, and tunnel identifiers:

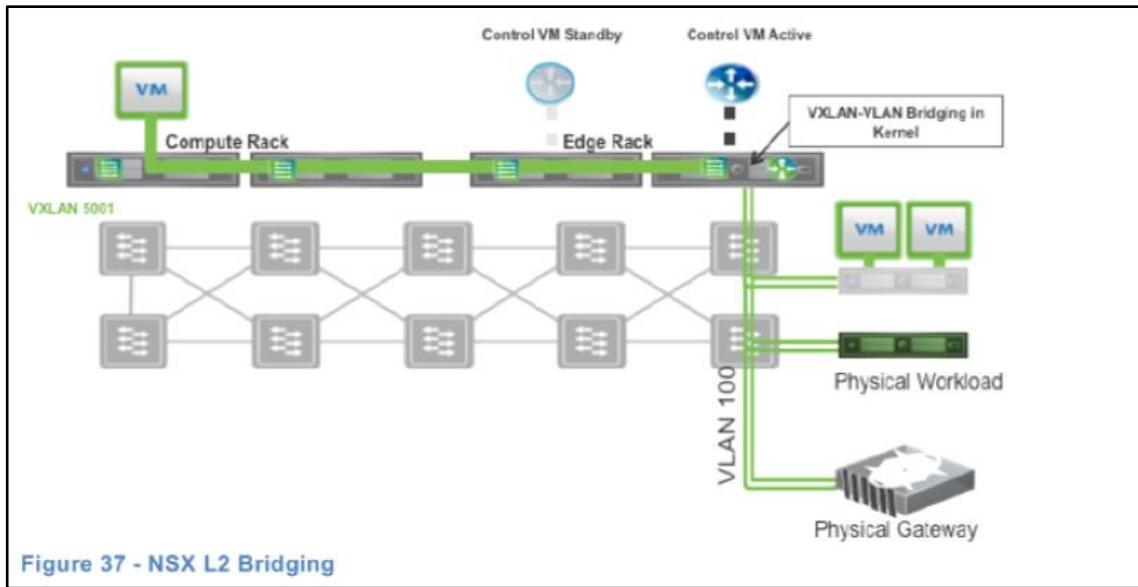


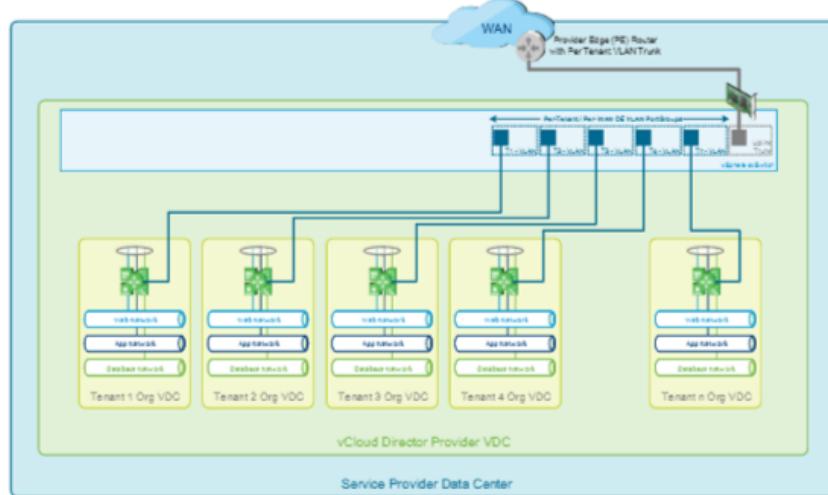
Figure 37 - NSX L2 Bridging

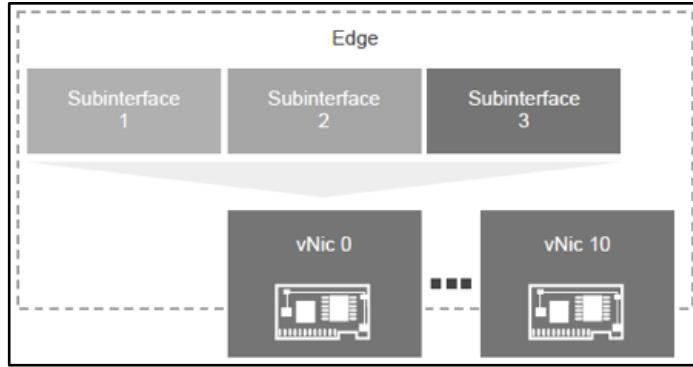
Figure 37 shows an example of L2 bridging. A VM connected in logical space to the VXLAN segment 5001 needs to communicate with a physical device deployed in the same IP subnet but connected to a physical network infrastructure through VLAN 100. The VXLAN-VLAN bridging configuration is part of the distributed router configuration; it is running on the same ESXi host where the control VM is located. This control VM serves as an anchor for the location of

When the first VM connects to a VXLAN segment – VM1 on ESXi1 and VM2 on ESXi-2 in this example – the ESXi host generates a control plane message to the specific controller node in charge of that specific logical switch slice with VNI/VTEP mapping information. The controller node populates its local VTEP table with this information and sends a report message to all ESXi hypervisors hosting VMs actively connected to that same VXLAN segment. In this example, The final piece of information shared with the controller is the IP address of the VMs. This controller populates its local ARP table in order to perform the ARP suppression functionality discussed the “Unicast Traffic (Virtual to Virtual Communication)” section.

109. The exemplary NSX system described above can receive, over a tunnel, a transmission on a physical interface having an interface identifier, the transmission identifying a tunnel identifier. For instance, an NSX Edge node can receive client transmissions sent through a tunnel (e.g., an 802.1q tunnel) via a physical port (e.g., an NSX Edge I/O port), such that the transmission includes a VLAN ID and physical port number:

Figure 17. Trunking Multiple External Networks to a vCloud Director Environment





110. The exemplary NSX system described above can further determine correct customer forwarding information from the customer lookup information using the physical interface identifier and the tunnel identifier. For example, the customer forwarding information, which includes the private destination virtual server address, is determined using the associated incoming physical port and VLAN ID:

Multitenant Data Center Networking in vSphere

Each customer has a separate WAN CE router. Because the connection to their tenant environments could, therefore have overlapping addresses (from within their vCloud Director organization, or from their WAN), each must be separated through the data center and into the vCloud Director managed environment. This typically means that each tenant's WAN connection is presented as a separate external network with a separate VLAN ID, and therefore requires a separate VLAN-backed port group in the vSphere dvSwitch to connect to the WAN interface of their respective Edge Services Gateways.

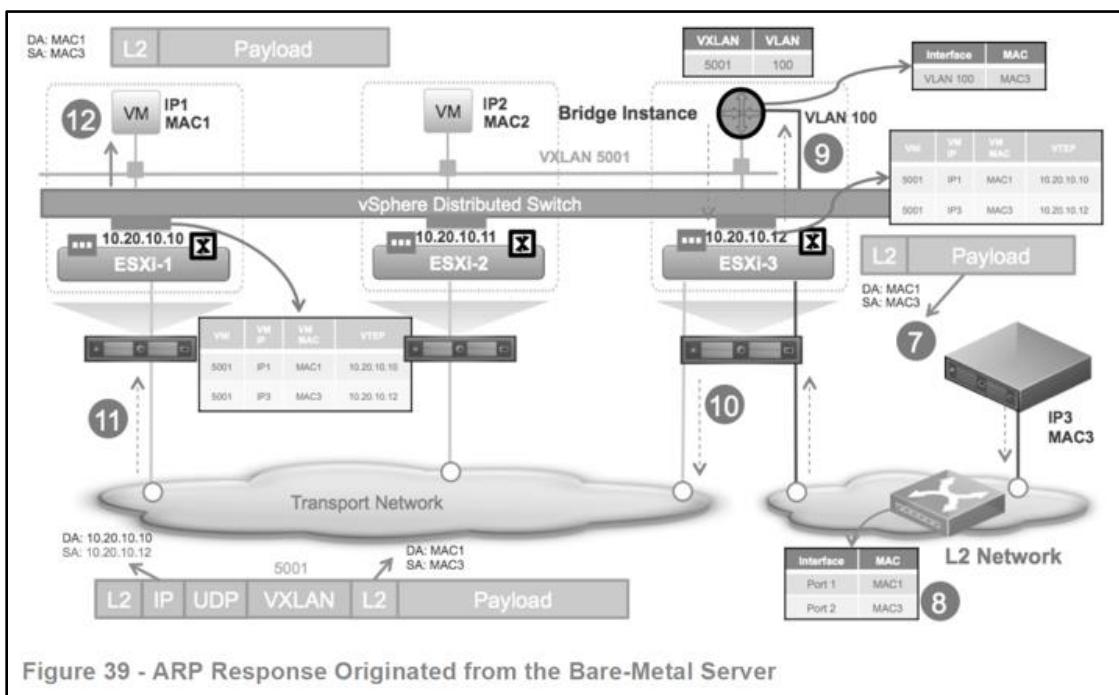


Additional deployment considerations for the NSX L2 bridging include:

The VXLAN-VLAN mapping is always performed in a 1:1 fashion. Traffic for a given VXLAN can only be bridged to a specific VLAN, and vice versa. A given bridge instance for a specific VXLAN-VLAN pair is always active on a single ESXi host.

The Tunnel ID is a construct that is used to map/associate the networks between sites. In Figure 2 and Figure 3, the Tunnel ID 1203 maps the VXLAN 6002 on the L2VPN server side (VMware Cloud Provider) of the VPN tunnel to VLAN 203 on the L2VPN client side (customer) of the VPN tunnel.

111. The exemplary NSX system described above can use the customer forwarding information to identify a physical interface and tunnel identifier associated with a network address of the transmission. For example, using the customer forwarding information NSX can determine the VXLAN ID, VTEP, and outgoing physical port associated with the private IP address of the target virtual server:



4.2.2 Populating the Controller Tables

Controller tables handle information essential for L2 unicast communication. Control plane communication between ESXi hosts and the controller cluster is used to populate the VTEP, MAC, and ARP tables on controller nodes. This process is detailed in Figure 32.

112. The exemplary NSX system described above can send the transmission to the network address on the identified physical interface using the identified tunnel identifier. For instance, NSX sends the traffic to the determined virtual server's private IP address by routing the traffic through the identified physical port and tunnel (e.g. VXLAN tunnel) associated with the virtual server as indicated by the customer forwarding information:

When the first VM connects to a VXLAN segment – VM1 on ESXi1 and VM2 on ESXi-2 in this example – the ESXi host generates a control plane message to the specific controller node in charge of that specific logical switch slice with VNI/VTEP mapping information. The controller node populates its local VTEP table with this information and sends a report message to all ESXi hypervisors hosting VMs actively connected to that same VXLAN segment. In this example, the final piece of information shared with the controller is the IP address of the VMs. This controller populates its local ARP table in order to perform the ARP suppression functionality discussed in the "Unicast Traffic (Virtual to Virtual Communication)" section.

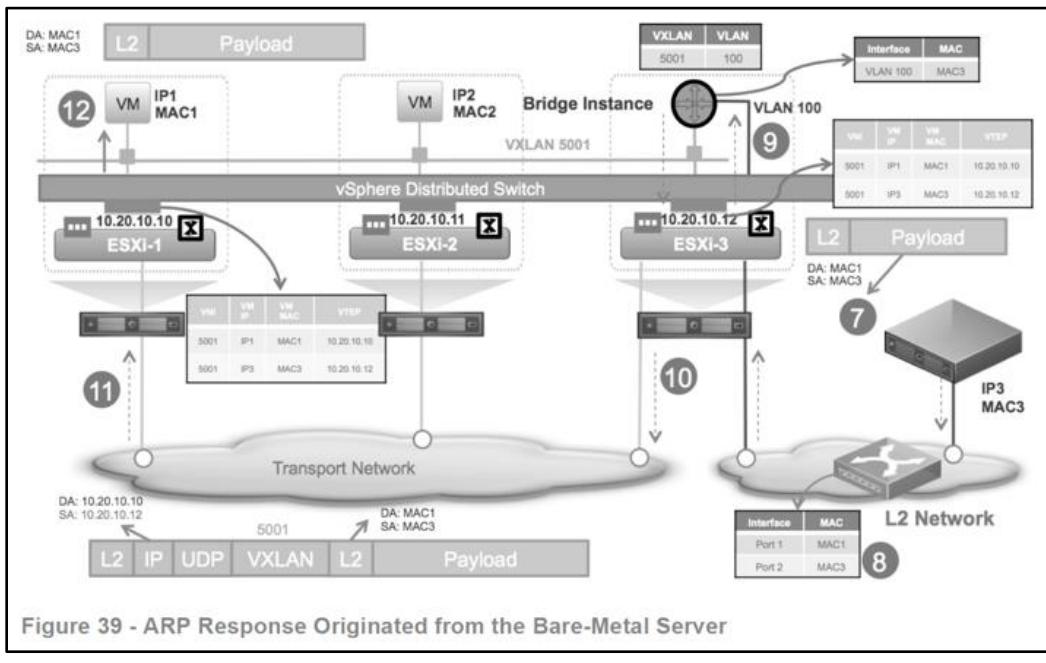


Figure 39 - ARP Response Originated from the Bare-Metal Server

113. In order to obtain the benefit of the NSX multi-tenant functionality, VMware's customers must operate the NSX system described above as required by VMWare, and therefore, if any action by VMware's customer is necessary for direct infringement of at least claims 1 and 3-6, such actions are attributable to VMware.

114. Additionally, VMware has been, and currently is, an active inducer of infringement of the RE '051 patent under 35 U.S.C. § 271(b) and contributory infringer of the RE '051 patent under 35 U.S.C. § 271(c) either literally and/or by the doctrine of equivalents.

115. VMware has actively induced, and continues to actively induce, infringement of the RE '051 patent by intending that others use, offer for sale, or sell in the United States, products and/or services covered by one or more claims of the RE '051 patent, including but not limited to VMware's NSX, NSX-T, vCloud Director, vSphere, ESXi, and any other VMware products and/or services, either alone or in combination, that operate in materially the same manner. VMware provides these products and/or services to others, such as customers, resellers and end-user customers, who, in turn, use, provision for use, offer for sale, or sell in the United States products and/or services that directly infringe one or more claims of the RE '051 patent.

116. VMware has contributed, and continues to contribute to, the infringement of the RE '051 patent by others by knowingly providing products and/or services that when installed and configured result in a system as intended by VMware, that directly infringes one or more claims of the RE '051 patent.

117. VMware knew of the RE '051 patent, or should have known of the RE '051 patent, but was willfully blind to its existence. Upon information and belief, VMware has had actual knowledge of the RE '051 patent since at least as early as the service upon VMware of this Complaint. By the time of trial, VMware will have known and intended (since receiving such

notice) that its continued actions would infringe and actively induce and contribute to the infringement of one or more claims of the RE ‘051 patent.

118. VMware has committed, and continues to commit, affirmative acts that cause infringement of one or more claims of the RE ‘051 patent with knowledge of the RE ‘051 patent and knowledge or willful blindness that the induced acts constitute infringement of one or more claims of the RE ‘051 patent. As an illustrative example only, VMware induces such acts of infringement by its affirmative actions of intentionally providing hardware and/or software components that when used in their normal and customary way as designed and intended by VMware, infringe one or more claims of the RE ‘051 patent and/or by directly or indirectly providing instructions on how to use its products and services in a manner or configuration that infringes one or more claims of the RE ‘051 patent, including those found at one or more of the following:

- www.vmware.com/products/horizon.html
- www.vmware.com/products/esxi-and-esx.html
- www.vmware.com/products/vsphere.html
- www.vmware.com/products/nsx.html
- docs.vmware.com/en/VMware-NSX-T-Data-Center/index.html
- www.vmware.com/products/vcloud-director.html
- www.vmware.com/files/pdf/vmware-it-journey/vmware-nsx-in-multitenant-cloud-cs.pdf
- www.youtube.com/watch?v=3OvrKZYnjM
- www.youtube.com/watch?v=AUbMqcIAg1g

- www.youtube.com/watch?v=EvXn2QiL3gs
- www.youtube.com/watch?v=xpYDM6sZOWY
- www.youtube.com/watch?v=D3pOo1I1rm0

119. VMware has also committed contributory infringement by, *inter alia*, knowingly selling products and/or services that when used cause the direct infringement of one or more claims of the RE ‘051 patent by a third party, and which have no substantial non-infringing uses, or include a separate and distinct component that is especially made or especially adapted for use in infringement of the RE ‘051 patent and is not a staple article or commodity of commerce suitable for substantial non-infringing use.

120. Upon information and belief, VMware has known of the RE ‘051 patent since at least December 13, 2013 when VMware’s 8,619,771 patent issued citing the RE ‘051 as a prior art reference.

121. As a result of VMware’s acts of infringement, Plaintiffs have suffered and will continue to suffer damages in an amount to be proved at trial.

COUNT V
(VMware’s Infringement of U.S. Patent No. RE 44,818)

122. Paragraphs 1- 121 are reincorporated by reference as if fully set forth herein.

123. The elements claimed by the RE ‘818 patent, taken alone or in combination, were not well-understood, routine or conventional to one of ordinary skill in the art at the time of the invention of the RE ‘818 patent, around 2007. Rather, the patent teaches and claims an improved way to facilitate and manage virtualized input/output (“I/O”) subsystems (e.g. virtual I/O servers) that represented a novel and non-obvious approach not present in the state of the art at that time. The invention improved upon then-existing virtual I/O servers, which are continuously engaged in

network communications, by enabling more granular QoS (Quality of Service) controls to be applied to those communications. The invention improved prior art virtual I/O systems by using a hierarchical token bucket allocator to determine how networking resources are allocated to communications being sent by virtual I/O servers over a local area network (“LAN”) connecting the virtual I/O server to a physical storage resource such as a hard drive. Specifically, the invention requires maintaining a connection over a network fabric from a virtual machine to a virtual network interface layer of the virtual I/O server, sending communications (i.e., LAN packets) relating to I/O operations from the VM to the virtual I/O server over the network fabric, and, prior to transmitting the LAN packets over the local area network, enforcing I/O bandwidth limitations via a hierarchical token bucket resource allocation such that the I/O packets are only transmitted across the LAN if there is a sufficient amount of remaining tokens.

124. This is to be contrasted with then-existing systems that for example, did not enforce more granular QoS (Quality of Service) requirements on virtual servers as they engaged in I/O operations over a LAN. This resulted in an inefficient allocation of network resources to such virtual I/O server communications, which in turn lead to suboptimal virtual I/O server performance. Virtual I/O servers, for example, were not previously able to obtain tailored LAN or SANaccess allocations, that are today required to handle the wide variety of communications typical of virtual I/O servers. Examples of such typical virtual I/O communications range from high bandwidth, latency-tolerant communications such as virtual storage file transfers on one hand, to low bandwidth, low-latency, bursty communication such as VoIP (“voice-over-IP”) telephone calls on the other hand.

125. The invention represented a technical solution to an unsolved technological problem. The written description of the RE ‘818 patent describes, in technical detail, each of the

limitations in the claims, allowing a person of skill in the art to understand what those limitations cover, and therefore what was claimed, and also understand how the non-conventional and non-generic ordered combination of the elements of the claims differ markedly from what had been conventional or generic in the industry at the time of the invention of the RE ‘818 patent. More specifically, the claims of the RE ‘818 patent require maintaining a connection over a network fabric to a virtual network interface layer of an application server, presenting at a physical network interface a virtual node identifier to a LAN, enforcing a hierarchical token bucket resource allocation of bandwidth across the LAN network interface, receiving over the network fabric connection LAN packets destined for a target reachable through the LAN, classifying packets relative to the hierarchical token bucket resource allocation to determine the amount of tokens available, comparing the size of the received packets to the amount of available tokens, forwarding the received packets across the LAN to and from the target if the amount of available tokens is sufficient, and buffering the packets if the amount of tokens is insufficient.

126. The systems and methods covered by the asserted claims therefore differ markedly from the conventional and generic systems in use at the time of this invention, which *inter alia* lacked the ability to enforce bandwidth constraints on the I/O traffic to and from the virtual servers using hierarchical token bucket resource allocation. Such functionality in turn enables resource allocation decisions that meet QoS (Quality of Service) requirements to be made at multiple points/layers in the virtual computing system (e.g., LAN allocation decisions can be made at a server level with respect to a type of traffic, or at a physical port level or virtual port level with respect to traffic moving to/from a specific server).

127. As described above, the RE ‘818 patent is drawn to solving a specific, technical problem arising in the context of facilitating and managing I/O operations over a LAN in a

virtualized environment. Consistent with the problem addressed being rooted in such management techniques in a virtualized computing environment over a LAN, the RE ‘818 patent’s solutions naturally are also rooted in that same technology and cannot be performed with pen and paper or in the human mind.

128. VMware has directly infringed, and continues to directly infringe, literally and/or by the doctrine of equivalents, individually and/or jointly, the RE ‘818 patent by making, using, testing, selling, offering for sale and/or importing into the United States products and/or services covered by at least claims 1, 17, 30, 32-33 and 37-42 of the RE ‘818 patent. VMware’s products and/or services that infringe the RE ‘818 patent include, but are not limited to, vSphere, vCenter Server, VMware’s ESX/ESXi hypervisor, vSAN and virtual distributed switch (“VDS”) functionality, and any other VMware products and/or services, either alone or in combination, that operate in materially the same manner.

129. Claim 42 of the RE ‘818 patent is reproduced below:

42. A method of facilitating management of input/output subsystems in a virtual input/output server, the method comprising:

maintaining a connection, over a network fabric, to a virtual network interface layer of an application server, to receive local area network packets;

presenting, at a physical network interface, a virtual node identifier to a local area network;

enforcing a hierarchical token bucket resource allocation of bandwidth across the physical network interface;

receiving, over the connection, local area network packets destined for a target on the local area network, thereby resulting in received local area network packets;

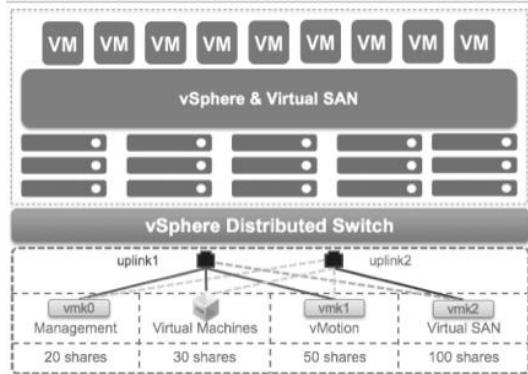
classifying the received local area network packets relative to the hierarchical token bucket resource allocation to determine a current amount of tokens available;

comparing a size of the received local area network packets to the current amount of tokens available;

forwarding the received local area packets across the physical local area network to and from the target, if the current amount of tokens available are sufficient; and

buffering the received local area network packets, if the current amount of tokens available are insufficient.

130. As one non-limiting example, a deployment of vSphere, vSAN, VDS, and ESX/ESXi provides for facilitating management of input/output subsystems in a virtual input/output server. A vSphere deployment supports a VM-based application server that utilizes a shared vSAN virtual storage system and VDS. VMs in a vSphere deployment can perform a number of I/O operations within the virtualized environment, including virtual storage operations as discussed in the example highlighted below:



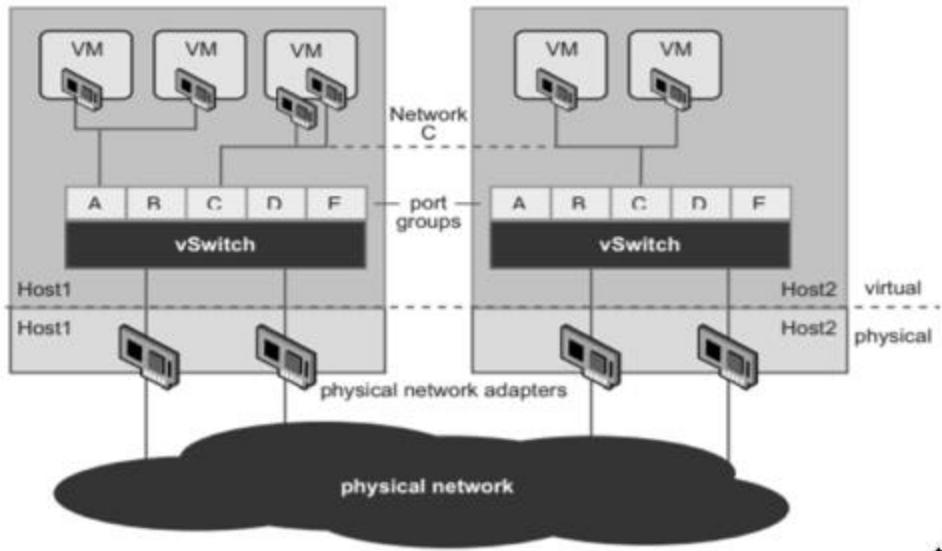
VSAN Networking

Network connectivity is the heart of any VSAN cluster. VSAN cluster hosts use the network for virtual machine (VM) I/O and also communicate their state between one another. Consistent and correct network configuration is key to a successful VSAN deployment. Because the majority of disk I/O will either come from a remote host, or will need to go to a remote host, VMware recommends leveraging a 10 GbE infrastructure. Note that although 1 GbE is fully supported in hybrid configurations, it could become a bottleneck in large-scale deployments.

VMware vSphere provides two different types of virtual switch, both of which are fully supported with VSAN:

- The VMware standard virtual switch (VSS) provides connectivity from VMs and VMkernel ports to external networks but is local to an ESXi host.
- A vSphere Distributed Switch (VDS) gives central control of virtual switch administration across multiple ESXi hosts. A VDS can also provide additional networking features over and above what a VSS can offer, such as network I/O control (NIOC) that can provide quality of service (QoS) on your network. Although a VDS normally requires a particular vSphere edition, VSAN includes a VDS independent license of the vSphere edition you are running.

131. Continuing with the above non-limiting example, the vSphere deployment including vSAN and VDS can maintain a connection, over a network fabric, to a virtual network interface layer of an application server to receive LAN packets. Each VM presents a vNIC (which corresponds to a NIC or Network Interface Card of the physical host) to the VDS through which the VDS can receive I/O traffic, via a LAN, from the VM:



132. Furthermore, a vSphere deployment with vSAN and VDS can present, at a physical network interface, a virtual node identifier to a LAN. For example, when a VM generates a vSAN I/O operation the request traverses the VM's virtual ports to the host's physical NIC and onto the LAN connecting the physical hosts comprising the vSAN. Each physical host is associated with at least one virtual port/port group, e.g., a VMkernel port, which has a unique ID and is used to transport vSAN I/O operations:

VMkernel Network

On each ESXi host that wants to participate in a VSAN cluster, a VMkernel port for VSAN communication must be created. The VMkernel port is labeled Virtual SAN traffic and was introduced in vSphere 5.5. This VMkernel port is used for intra-cluster node communication. It is also used for reads and writes when one of the ESXi hosts in the cluster owns a particular VM, but the actual data blocks that make up the VM files are located on a different ESXi host in the cluster. In this case, I/O will need to traverse the network configured between the hosts in the cluster, as depicted in [Figure 2.4](#), where VMkernel interface vmk2 is used for VSAN traffic by all the hosts in the VSAN cluster. The VM residing on ESXi-01 does all of its reads and writes leveraging the VSAN network.

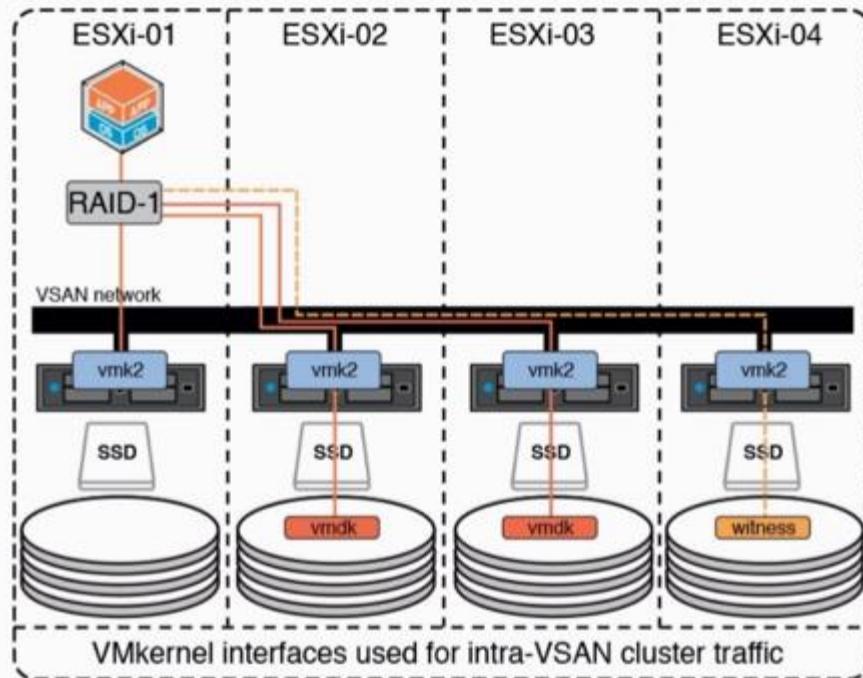


Figure 2.4 VSAN traffic

133. Continuing with the above example, the vSphere deployment with vSAN and VDS enforces a hierarchical token bucket resource allocation of bandwidth across physical network interfaces in a vSphere environment. This enforcement can be applied at varying levels of abstraction within the virtualized environment, e.g., at the VDS level, virtual port level and/or vNIC level, thus achieving a hierarchical bandwidth allocation. For instance, vSAN I/O operations can be allocated an overall amount of bandwidth “shares” (i.e. tokens) at the VDS level, and each member VM within that VDS can further be allocated a sub-portion of those available shares:

How Network I/O Control Works

Network I/O Control enforces the share value specified for the different traffic types only when there is network contention. When contention occurs Network I/O Control applies the share values set to each traffic type. As a result, less important traffic, as defined by the share percentage, will be throttled, allowing more important traffic types to gain access to more network resources.

Network I/O Control also allows the reservation of bandwidth for system traffic based on the capacity of the physical adapters on a host, and enables fine-grained resource control at the virtual machine network adapter level. Resource control is similar to the model for vCenter CPU

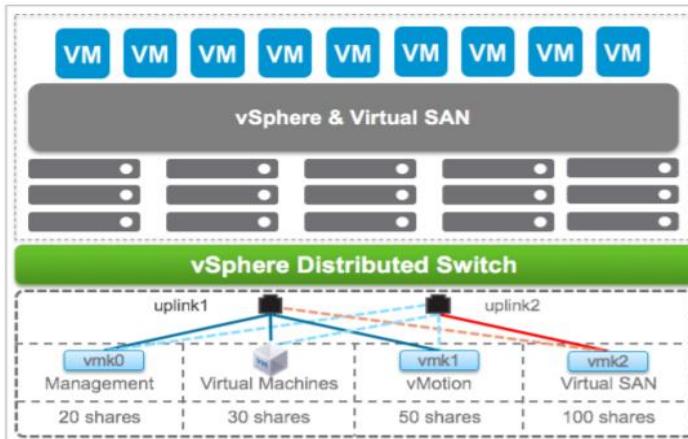


Figure 3: VMware Distributed Switched - QoS with Network I/O Control

134. The vSphere deployment with vSAN and VDS further can receive, over the connection, LAN packets destined for a target on the LAN, classify the received LAN packets relative to the hierarchical token bucket resource allocation to determine a current amount of tokens available, and compare a size of the received LAN packets to the current amount of tokens available. For example, the vSAN I/O operation packets are sent via the VMkernal port to the physical network interface destined for a vSAN network element accessible through the LAN. The size of the packets is calculated and the system compares the size to the amount of remaining ‘shares’ allocated for vSAN traffic bandwidth:

4.1 Network I/O Control

vSphere Network I/O Control, a feature available with vSphere Distributed Switches, is a mechanism to implement a Quality of Service (QoS) on network traffic. This can be extremely useful for vSAN when vSAN traffic does not have its own dedicated network interface card (NIC) and has to share the physical NIC with other traffic types, e.g. vMotion, Management, virtual machines.

4.3 Reservation, Shares and Limits

Setting **shares** makes a certain bandwidth available to vSAN when the physical adapter assigned for vSAN becomes saturated. This prevents vSAN from consuming the entire capacity of the physical adapter during rebuild and synchronization operations. For example, the physical adapter might become saturated when another physical adapter in the team fails and all traffic in the port group is transferred to the remaining adapter(s) in the team. The shares mechanism ensures that no other traffic impact the v SAN network, and vice versa.

4.5 NIOC Configuration Example

Traffic Type	Shares	Value
vSAN	High	100
vSphere vMotion	Low	25
Virtual machine	Normal	50
iSCSI/NFS	Low	25

If the 10-GbE adapter becomes saturated, Network I/O Control allocates 5Gbps to vSAN on the physical adapter, 3.5Gbps to vMotion and 1.5Gbp to virtual machine traffic. The above values may be

135. Continuing with the above example, the vSphere deployment with vSAN and VDS forwards the received LAN packets across the physical LAN to and from the target, if the current amount of tokens available are sufficient, and buffers the received packets if the current amount of tokens available are insufficient. For instance, when the system determines that there is sufficient remaining ‘shares’ allocated for vSAN traffic bandwidth at each applicable hierachal level, the vSAN I/O packets are sent through the physical LAN to the target vSAN datastore. If there is insufficient remaining vSAN bandwidth shares to transport the packet, the packet is buffered until sufficient shares become available and then forwarded:

In this blog post, we go into the trenches of the (Distributed) vSwitch with a focus on vSphere ESXi network IOChain. It is important to understand the core constructs of the vSphere networking layers for i.e. troubleshooting connectivity issues. In a second blog post on this topic, we will look closer into virtual network troubleshoot tooling.

Uplink level

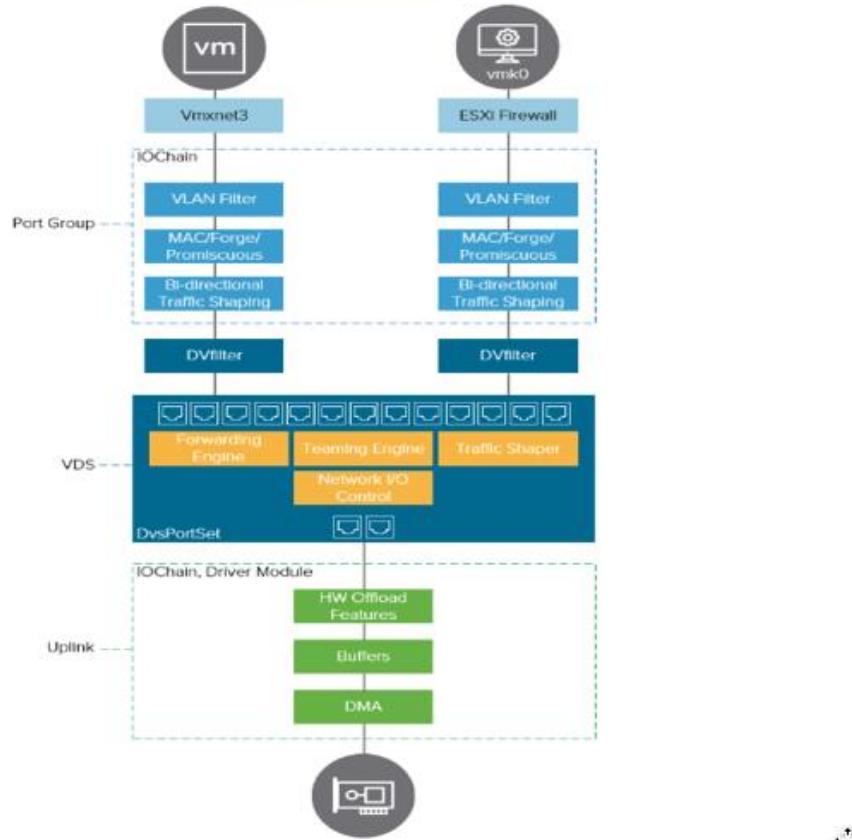
At this level, the traffic sent from the vSwitch to an external host finds its way to the driver module. This is where all the hardware offloading is taking place. The Supported hardware offloading features depends strongly on the physical NIC in combination with a specific driver module. Typically supported hardware offloading functions that in NICs are TCP Segment Offload (TSO), Large Receive Offload (LRO) or Checksum Offload (CSO). Network overlay protocol offloading like with VXLAN and Geneve, as used in NSX-v and NSX-T respectively, are widely supported on modern NICs.

Next to hardware offloading, the buffer mechanisms come into play in the Uplink level. I.e., when processing a burst of network packets, ring buffers come into play. Finally, the bits transmit onto the DMA controller to be handled by the CPU and physical NIC onwards to the Ethernet fabric.

Distributed Switch

The magic happens when vSphere VDS lets us span network switch configurations over multiple ESXi hosts that can be grouped over multiple clusters even. The VDS is available in the Enterprise licensing model or shipped with your vSAN license. It allows network information to be distributed over all member ESXi hosts, thus providing network configuration consistency in your vSphere environment.

An extensive feature-set is provided for quality control features like Network I/O Control (NIOC), in- and egress traffic shaping and traffic flow monitoring options. Next to that, it allows for additional teaming options like LBT (Load Balanced Teaming) and LACP (Link Aggregation Control Protocol). Review all features that are listed on the [vSphere Distributed Switch page](#).



136. VMware and its customers operating, for example, vSphere environments with vSAN and VDS functionality directly infringe at least claims 1, 17, 30, 32-33 and 37-42 of the ‘818 patent. In order to obtain the benefit of the vSphere deployment, including the vSAN and

VDS features, VMware's customers must operate vSphere as required by VMware and therefore, if any action by VMware's customer is necessary for direct infringement of any claims (including at least claims 1, 17, 30, 32-33 and 37-42), such actions are attributable to VMware.

137. Additionally, VMware has been, and currently is, an active inducer of infringement of the RE '818 patent under 35 U.S.C. § 271(b) and contributory infringer of the RE '818 patent under 35 U.S.C. § 271(c) either literally and/or by the doctrine of equivalents.

138. VMware has actively induced, and continues to actively induce, infringement of the RE '818 patent by intending that others use, offer for sale, or sell in the United States, products and/or services covered by one or more claims of the RE '818 patent, including but not limited to VMware's vSphere platform, vCenter Server, vSAN and VDS functionality, ESX/ESXi hypervisor, and any VMware products and/or services, either alone or in combination, that operate in materially the same manner. VMware provides these products and/or services to others, such as customers, resellers and end-user customers, who, in turn, use, provision for use, offer for sale, or sell in the United States products and/or services that directly infringe one or more asserted claims of the RE '818 patent.

139. VMware has contributed to, and continues to contribute to, the infringement of the RE '818 patent by others by knowingly providing products and/or services that, when installed and configured result in a system as intended by VMware, directly infringes one or more claims of the RE '818 patent.

140. VMware knew of the RE '818 patent, or should have known of the RE '818 patent, but was willfully blind to its existence. Upon information and belief, VMware has had actual knowledge of the RE '818 patent since at least as early as the service upon VMware of this Complaint. By the time of trial, VMware will have known and intended (since receiving such

notice) that its continued actions would infringe and actively induce and contribute to the infringement of one or more claims of the RE ‘818 patent.

141. VMware has committed, and continues to commit, affirmative acts that cause infringement of one or more claims of the RE ‘818 patent with knowledge of the RE ‘818 patent and knowledge or willful blindness that the induced acts constitute infringement of one or more claims of the RE ‘818 patent. As an illustrative example only, VMware induces such acts of infringement by its affirmative action by intentionally providing hardware and/or software components that when used in their normal and customary way as designed and intended by VMware, infringe one or more claims of the RE ‘818 patent and/or by directly or indirectly providing instructions on how to use its products and services in a manner or configuration that infringes one or more claims of the RE ‘818 patent, including those found at one or more of the following:

- <https://www.vmware.com/products/vsphere.html>
- <https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.vsan.doc/GUID-ACC10393-47F6-4C5A-85FC-88051C1806A0.html>
- <https://www.youtube.com/watch?v=UteOfQ85064&list=PLjwkgfjHppDux1XhPB8pW3vS43Aglfq2c&index=4>
- <https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/vsan/vmware-virtual-san-data-locality.pdf>
- <https://www.vmware.com/files/pdf/products/vsan/vmware-vs SAN-layer2-and-layer3-network-topologies.pdf>

- <https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/vsan/vmware-vs-san-66-licensing-guide.pdf>
- <https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.REvsphereREvirtualsan.doc/GUID-AFF133BC-F4B6-4753-815F-20D3D752D898.html>
- <https://docs.vmware.com/en/VMware-Validated-Design/4.0/com.vmware.vvd.mseg-design.doc/GUID-586A815C-86D3-4B86-BB8B-DCF5BEBE5EE7.html>
- <https://docs.vmware.com/en/VMware-vSphere/6.7/vsan-671-planning-deployment-guide.pdf>
- http://download3.vmware.com/vmworld/2014/downloads/session-pdfs/STO1279_Final_US.pdf
- https://storagehub.vmware.com/export_to_pdf/vmware-r-vs-san-tm-network-design

142. VMware has also committed, and continues to commit, contributory infringement by, *inter alia*, knowingly selling products and/or services that when used cause the direct infringement of one or more claims of the RE ‘818 patent by a third party, and which have no substantial non-infringing uses, or include a separate and distinct component that is especially made or especially adapted for use in infringement of the RE ‘818 patent and is not a staple article or commodity of commerce suitable for substantial non-infringing use.

143. As a result of VMware’s acts of infringement, Plaintiffs have suffered and will continue to suffer damages in an amount to be proved at trial.

PRAYER FOR RELIEF

Plaintiffs request that the Court enter judgment against VMware:

- (A) that VMware has infringed one or more claims of each of the above patents-in-suit, directly and/or indirectly, literally and/or under the doctrine of equivalents;
- (B) awarding damages sufficient to compensate Plaintiffs for VMware's infringement under 35 U.S.C. § 284;
- (C) finding this case exceptional under 35 U.S.C. § 285 and awarding Plaintiffs their reasonable attorneys' fees;
- (D) awarding Plaintiffs' their costs and expenses incurred in this action;
- (E) awarding Plaintiffs prejudgment and post-judgment interest; and
- (F) granting Plaintiffs such further relief as the Court deems just and appropriate.

DEMAND FOR JURY TRIAL

Plaintiffs demand trial by jury of all claims so triable under Federal Rule Of Civil Procedure 38.

Date: July 31, 2019

Respectfully submitted,

/s/ Derek Gilliland
DEREK GILLILAND
STATE BAR NO. 24007239
NIX PATTERSON L.L.P.
222 N. Fredonia
Longview, Texas 75601
903.215.8310 (telephone)
dgilliland@nixlaw.com

KARL RUPP
State Bar No. 24035243
NIX PATTERSON L.L.P.
1845 Woodall Rodgers Fwy., Suite 1050
Dallas, Texas 75201
972.831.1188 (telephone)
972.444.0716 (facsimile)
krupp@nixlaw.com

OF COUNSEL:

Paul J. Hayes
phayes@princelobel.com
Matthew Vella
mvella@princelobel.com
Robert R. Gilman
rgilman@princelobel.com
Jonathan DeBlois
jdeblois@princelobel.com
Alex Breger
abreger@princelobel.com
PRINCE LOBEL TYE LLP
One International Place, Suite 3700
Boston, MA 02110
Tel: (617) 456-8000
Fax: (617) 456-8100

COUNSEL for PLAINTIFFS